

2.6.10-rc1 & ahci & IHC6R & 925X & raid1

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2004-10/7212.html>

From: Matthijs Melchior (*mmelchior_at_xs4all.nl*)

Date: 10/23/04

Date: Sat, 23 Oct 2004 17:51:41 +0200
To: linux-kernel@vger.kernel.org

Hi,

Linux running on my new Dell-8400 computer does not recognize the SATA disks [No problem with the PATA drives].

The attached log is output from 'dmesg' with '/proc/modules' and 'ps -Hel' appended.

This machine has 4 disks, configured in the Intel Raid configuration BIOS as 3 devices, with disks on ports 0 & 2 being a raid1 pair, half of disks 1 & 3 also a raid1 device and the rest a raid0 device. The raid signatures on disk 3 have been removed [for testing].

The "SATA Operation" setting in BIOS is set to "RAID Autodetect / AHCI".

I have set symbols ATA_DEBUG and ATA_VERBOSE_DEBUG to 1 for more output and built a new set of modules.

I have configured netbooting [PXE] for this machine and I can boot the kernel with an initrd. The only thing it currently does is loading usb modules for console, usb keyboard, network and starting a shell.

Giving the command "modprobe ahci" hangs the machine.

Using "modprobe ahci &" instead and investigating the state at that time leads me to suspect ata_dev_identify() waits for an interrupt that is not occurring.

The attached '/proc/modules' and 'ps -Hel' are from this environment.

Please, can someone help me debugging this further, what other info would be usefull, are there other experiments that can provide new insights into this problem.....

--

Regards,

Matthijs Melchior
mmelchior@xs4all.nl

----- -o)
Maarssen /\\
Netherlands __

Linux version 2.6.10-rc1 (mmelchio@p2pc) (gcc version 3.2.2) #1 SMP Sat Oct 23 06:23:17 CEST 2004

BIOS-provided physical RAM map:

BIOS-e820: 0000000000000000 - 00000000000a0000 (usable)
BIOS-e820: 00000000000f0000 - 0000000000100000 (reserved)
BIOS-e820: 0000000000100000 - 0000000003fe8cc00 (usable)
BIOS-e820: 0000000003fe8cc00 - 0000000003fe8ec00 (ACPI NVS)
BIOS-e820: 0000000003fe8ec00 - 0000000003fe90c00 (ACPI data)
BIOS-e820: 0000000003fe90c00 - 0000000040000000 (reserved)
BIOS-e820: 00000000e0000000 - 00000000f0000000 (reserved)
BIOS-e820: 00000000fec00000 - 00000000fed00400 (reserved)
BIOS-e820: 00000000fed20000 - 00000000feda0000 (reserved)
BIOS-e820: 00000000fee00000 - 00000000fef00000 (reserved)
BIOS-e820: 00000000ffb00000 - 0000000100000000 (reserved)

126MB HIGHMEM available.

896MB LOWMEM available.

found SMP MP-table at 000fe710

On node 0 totalpages: 261772

DMA zone: 4096 pages, LIFO batch:1

Normal zone: 225280 pages, LIFO batch:16

HighMem zone: 32396 pages, LIFO batch:7

DMI 2.3 present.

ACPI: RSDP (v000 DELL) @ 0x000febf0

ACPI: RSDT (v001 DELL 8400 0x00000007 ASL 0x00000061) @ 0x000fccbc

ACPI: FADT (v001 DELL 8400 0x00000007 ASL 0x00000061) @ 0x000fccf8

ACPI: SSDT (v001 DELL st_ex 0x00001000 MSFT 0x0100000d) @ 0xfffc9180

ACPI: MADT (v001 DELL 8400 0x00000007 ASL 0x00000061) @ 0x000fcd6c

ACPI: BOOT (v001 DELL 8400 0x00000007 ASL 0x00000061) @ 0x000fcdde

ACPI: MCFG (v001 DELL 8400 0x00000007 ASL 0x00000061) @ 0x000fce06

ACPI: HPET (v001 DELL 8400 0x00000007 ASL 0x00000061) @ 0x000fce44

ACPI: DSDT (v001 DELL dt_ex 0x00001000 MSFT 0x0100000d) @ 0x00000000

ACPI: PM-Timer IO Port: 0x808

ACPI: Local APIC address 0xfe00000

ACPI: LAPIC (acpi_id[0x01] lapic_id[0x00] enabled)

Processor #0 15:3 APIC version 20

ACPI: LAPIC (acpi_id[0x02] lapic_id[0x01] enabled)

Processor #1 15:3 APIC version 20

ACPI: LAPIC (acpi_id[0x03] lapic_id[0x01] disabled)

ACPI: LAPIC (acpi_id[0x04] lapic_id[0x07] disabled)

ACPI: LAPIC_NMI (acpi_id[0xff] high level lint[0x1])

ACPI: IOAPIC (id[0x08] address[0xfec00000] gsi_base[0])

IOAPIC[0]: apic_id 8, version 32, address 0xfec00000, GSI 0-23

ACPI: INT_SRC_OVR (bus 0 bus_irq 0 global_irq 2 dfl dfl)

ACPI: INT_SRC_OVR (bus 0 bus_irq 9 global_irq 9 high level)

ACPI: IRQ0 used by override.

ACPI: IRQ2 used by override.

ACPI: IRQ9 used by override.

Enabling APIC mode: Flat. Using 1 I/O APICs

ACPI: HPET id: 0x8086a201 base: 0xfed00000

Using ACPI (MADT) for SMP configuration information

Built 1 zonelists

Kernel command line: ramdisk_size=16384 initrd=initrd root=/dev/ram0 vga=031B BOOT_IMAGE=linux

Initializing CPU#0

PID hash table entries: 4096 (order: 12, 65536 bytes)

Console: colour dummy device 80x25

Dentry cache hash table entries: 131072 (order: 7, 524288 bytes)

Inode–cache hash table entries: 65536 (order: 6, 262144 bytes)

Memory: 1029368k/1047088k available (1683k kernel code, 17116k reserved, 689k data, 188k init, 129584k highmem)

Checking if this processor honours the WP bit even in supervisor mode... Ok.

hpet0: at MMIO 0xf8800000, IRQs 2, 8, 0

hpet0: 0ns tick, 3 64–bit timers

Using HPET for base–timer

Using HPET for gettimeofday

Detected 2992.818 MHz processor.

Using hpet for high–res timesource

Calibrating delay loop... 5931.00 BogoMIPS (lpj=2965504)

Security Framework v1.0.0 initialized

Mount–cache hash table entries: 512 (order: 0, 4096 bytes)

CPU: After generic identify, caps: bfebfbff 00000000 00000000 00000000

CPU: After vendor identify, caps: bfebfbff 00000000 00000000 00000000

monitor/mwait feature present.

using mwait in idle threads.

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

CPU: L2 cache: 1024K

CPU: Physical Processor ID: 0

CPU: After all inits, caps: bfebfbff 00000000 00000000 00000080

Intel machine check architecture supported.

Intel machine check reporting enabled on CPU#0.

CPU0: Intel P4/Xeon Extended MCE MSRs (12) available

CPU0: Thermal monitoring enabled

Enabling fast FPU save and restore... done.

Enabling unmasked SIMD FPU exception support... done.

Checking 'hlt' instruction... OK.

CPU0: Intel(R) Pentium(R) 4 CPU 3.00GHz stepping 04

per–CPU timeslice cutoff: 2925.83 usecs.

task migration cache decay timeout: 3 msecs.

Booting processor 1/1 eip 3000

Initializing CPU#1

Calibrating delay loop... 5980.16 BogoMIPS (lpj=2990080)

CPU: After generic identify, caps: bfebfbff 00000000 00000000 00000000

CPU: After vendor identify, caps: bfebfbff 00000000 00000000 00000000

monitor/mwait feature present.

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

CPU: L2 cache: 1024K

CPU: Physical Processor ID: 0

CPU: After all inits, caps: bfebfbff 00000000 00000000 00000080

Intel machine check architecture supported.

Intel machine check reporting enabled on CPU#1.

CPU1: Intel P4/Xeon Extended MCE MSRs (12) available

CPU1: Thermal monitoring enabled

CPU1: Intel(R) Pentium(R) 4 CPU 3.00GHz stepping 04
Total of 2 processors activated (11911.16 BogoMIPS).
ENABLING IO-APIC IRQs
..TIMER: vector=0x31 pin1=2 pin2=-1
checking TSC synchronization across 2 CPUs: passed.
Brought up 2 CPUs
CPU0:
 domain 0: span 03
 groups: 01 02
 domain 1: span 03
 groups: 03
CPU1:
 domain 0: span 03
 groups: 02 01
 domain 1: span 03
 groups: 03
checking if image is initramfs...it isn't (ungzip failed); looks like an initrd
Freeing initrd memory: 4984k freed
NET: Registered protocol family 16
PCI: PCI BIOS revision 2.10 entry at 0xfb768, last bus=4
PCI: Using MMCONFIG
mtrr: v2.0 (20020519)
ACPI: Subsystem revision 20040816
ACPI: Interpreter enabled
ACPI: Using IOAPIC for interrupt routing
ACPI: PCI Root Bridge [PCI0] (00:00)
PCI: Probing PCI hardware (bus 00)
PCI: Ignoring BAR0-3 of IDE controller 0000:00:1f.1
PCI: Transparent bridge - 0000:00:1e.0
ACPI: PCI Interrupt Routing Table [_SB_.PCI0._PRT]
ACPI: PCI Interrupt Routing Table [_SB_.PCI0.PCI1._PRT]
ACPI: PCI Interrupt Routing Table [_SB_.PCI0.PCI2._PRT]
ACPI: PCI Interrupt Routing Table [_SB_.PCI0.PCI3._PRT]
ACPI: PCI Interrupt Routing Table [_SB_.PCI0.PCI4._PRT]
ACPI: PCI Interrupt Link [LNKA] (IRQs 3 4 5 6 7 9 10 *11 12 15)
ACPI: PCI Interrupt Link [LNKB] (IRQs 3 4 5 6 7 9 *10 11 12 15)
ACPI: PCI Interrupt Link [LNKC] (IRQs *3 4 5 6 7 9 10 11 12 15)
ACPI: PCI Interrupt Link [LNKD] (IRQs 3 4 5 6 7 9 10 11 12 15) *0, disabled.
ACPI: PCI Interrupt Link [LNKE] (IRQs 3 4 *5 6 7 9 10 11 12 15)
ACPI: PCI Interrupt Link [LNKF] (IRQs 3 4 5 6 7 *9 10 11 12 15)
ACPI: PCI Interrupt Link [LNKG] (IRQs 3 4 *5 6 7 9 10 11 12 15)
ACPI: PCI Interrupt Link [LNKH] (IRQs 3 4 5 6 7 9 *10 11 12 15)
Linux Plug and Play Support v0.97 (c) Adam Belay
PnPBIOS: Scanning system for PnP BIOS support...
PnPBIOS: Found PnP BIOS installation structure at 0xc00fe2d0
PnPBIOS: PnP BIOS version 1.0, entry 0xf0000:0xe2f4, dseg 0x40
PnPBIOS: 19 nodes reported by PnP BIOS; 19 recorded by driver
PCI: Using ACPI for IRQ routing
ACPI: PCI interrupt 0000:00:01.0[A] -> GSI 16 (level, low) -> IRQ 169
ACPI: PCI interrupt 0000:00:1c.0[A] -> GSI 16 (level, low) -> IRQ 169
ACPI: PCI interrupt 0000:00:1c.1[B] -> GSI 17 (level, low) -> IRQ 177

ACPI: PCI interrupt 0000:00:1d.0[A] -> GSI 21 (level, low) -> IRQ 185
ACPI: PCI interrupt 0000:00:1d.1[B] -> GSI 22 (level, low) -> IRQ 193
ACPI: PCI interrupt 0000:00:1d.2[C] -> GSI 18 (level, low) -> IRQ 201
ACPI: PCI interrupt 0000:00:1d.3[D] -> GSI 23 (level, low) -> IRQ 209
ACPI: PCI interrupt 0000:00:1d.7[A] -> GSI 21 (level, low) -> IRQ 185
ACPI: PCI interrupt 0000:00:1e.2[A] -> GSI 23 (level, low) -> IRQ 209
ACPI: PCI interrupt 0000:00:1f.1[A] -> GSI 16 (level, low) -> IRQ 169
ACPI: PCI interrupt 0000:00:1f.2[C] -> GSI 20 (level, low) -> IRQ 217
ACPI: PCI interrupt 0000:00:1f.3[B] -> GSI 17 (level, low) -> IRQ 177
ACPI: PCI interrupt 0000:01:00.0[A] -> GSI 16 (level, low) -> IRQ 169
ACPI: PCI interrupt 0000:02:00.0[A] -> GSI 16 (level, low) -> IRQ 169
pnp: 00:00: ioport range 0x800–0x8df could not be reserved
pnp: 00:00: ioport range 0xc00–0xc7f has been reserved
Simple Boot Flag value 0x87 read from CMOS RAM was invalid
Simple Boot Flag at 0x7a set to 0x1
highmem bounce pool size: 64 pages
VFS: Disk quotas dquot_6.5.1
Dquot–cache hash table entries: 1024 (order 0, 4096 bytes)
devfs: 2004–01–31 Richard Gooch (rgooch@atnf.csiro.au)
devfs: boot_options: 0x0
Initializing Cryptographic API
vesafb: framebuffer at 0xd0000000, mapped to 0xf8880000, using 7680k, total 16384k
vesafb: mode is 1280x1024x24, linelength=3840, pages=3
vesafb: protected mode interface info at c000:5890
vesafb: scrolling: redraw
vesafb: Truecolor: size=0:8:8:8, shift=0:16:8:0
fb0: VESA VGA frame buffer device
isapnp: Scanning for PnP cards...
isapnp: No Plug & Play device found
hpet_acpi_add: no address or irqs in _CRS
serio: i8042 AUX port at 0x60,0x64 irq 12
serio: i8042 KBD port at 0x60,0x64 irq 1
Serial: 8250/16550 driver \$Revision: 1.90 \$ 48 ports, IRQ sharing enabled
ttyS0 at I/O 0x3f8 (irq = 4) is a 16550A
ttyS0 at I/O 0x3f8 (irq = 4) is a 16550A
io scheduler noop registered
io scheduler anticipatory registered
io scheduler deadline registered
io scheduler cfq registered
RAMDISK driver initialized: 16 RAM disks of 16384K size 1024 blocksize
NET: Registered protocol family 2
IP: routing cache hash table of 8192 buckets, 64Kbytes
TCP: Hash tables configured (established 262144 bind 65536)
NET: Registered protocol family 8
NET: Registered protocol family 20
Starting balanced_irq
ACPI: (supports S0 S1 S3 S4 S5)
ACPI wakeup devices:
VBTN PCI0 PCI1 PCI2 PCI3 PCI4 KBD USB0 USB1 USB2 USB3
RAMDISK: cramfs filesystem found at block 0
RAMDISK: Loading 4984KiB [1 disk] into ram disk...


```
# modprobe ahci &
SCSI subsystem initialized
libata version 1.02 loaded.
ahci_init_one: ENTER
ahci version 0.11
ACPI: PCI interrupt 0000:00:1f.2[C] -> GSI 20 (level, low) -> IRQ 217
ahci_host_init: cap 0xc6107f03 port_map 0xf n_ports 4
ahci_host_init: mmio f8848c00 port_mmio f8848d00
ahci_setup_port: ENTER, base==0xf8848c00, port_idx 0
ahci_setup_port: base now==0xf8848d00
ahci_setup_port: EXIT
ahci_host_init: PORT_CMD 0x6
ahci_host_init: PORT_SCR_ERR 0x4050000
ahci_host_init: PORT_IRQ_STAT 0x0
ahci_host_init: mmio f8848c00 port_mmio f8848d80
ahci_setup_port: ENTER, base==0xf8848c00, port_idx 1
ahci_setup_port: base now==0xf8848d80
ahci_setup_port: EXIT
ahci_host_init: PORT_CMD 0x6
ahci_host_init: PORT_SCR_ERR 0x4050000
ahci_host_init: PORT_IRQ_STAT 0x0
ahci_host_init: mmio f8848c00 port_mmio f8848e00
ahci_setup_port: ENTER, base==0xf8848c00, port_idx 2
ahci_setup_port: base now==0xf8848e00
ahci_setup_port: EXIT
ahci_host_init: PORT_CMD 0x6
ahci_host_init: PORT_SCR_ERR 0x4050000
ahci_host_init: PORT_IRQ_STAT 0x0
ahci_host_init: mmio f8848c00 port_mmio f8848e80
ahci_setup_port: ENTER, base==0xf8848c00, port_idx 3
ahci_setup_port: base now==0xf8848e80
ahci_setup_port: EXIT
ahci_host_init: PORT_CMD 0x6
ahci_host_init: PORT_SCR_ERR 0x4050000
ahci_host_init: PORT_IRQ_STAT 0x0
ahci_host_init: HOST_CTL 0x80000000
ahci_host_init: HOST_CTL 0x80000002
PCI: Setting latency timer of device 0000:00:1f.2 to 64
ahci(0000:00:1f.2) AHCI 0001.0000 32 slots 4 ports 1.5 Gbps 0xf impl
ahci(0000:00:1f.2) flags: 64bit ncq pm led slum part
ata_device_add: ENTER
ata_host_add: ENTER
ata_port_start: prd alloc, virt f7d3d000, dma 37d3d000
ata1: SATA max UDMA/133 cmd 0xF8848D00 ctl 0x0 bmdma 0x0 irq 217
ata_host_add: ENTER
ata_port_start: prd alloc, virt f7e17000, dma 37e17000
ata2: SATA max UDMA/133 cmd 0xF8848D80 ctl 0x0 bmdma 0x0 irq 217
ata_host_add: ENTER
ata_port_start: prd alloc, virt f7d51000, dma 37d51000
ata3: SATA max UDMA/133 cmd 0xF8848E00 ctl 0x0 bmdma 0x0 irq 217
ata_host_add: ENTER
```

```

ata_port_start: prd alloc, virt c197b000, dma 197b000
ata4: SATA max UDMA/133 cmd 0xF8848E80 ctl 0x0 bmdma 0x0 irq 217
ata_device_add: probe begin
ata_device_add: ata1: probe begin
ahci_interrupt: ENTER
ahci_interrupt: port 0
ahci_interrupt: port 1
ahci_interrupt: port 2
ahci_interrupt: port 3
ahci_interrupt: EXIT
ahci_interrupt: ENTER
ahci_interrupt: port 0
ahci_interrupt: port 1
ahci_interrupt: port 2
ahci_interrupt: port 3
ahci_interrupt: EXIT
ata_dev_classify: found ATA device by sig
ata_dev_identify: ENTER, host 1, dev 0
ata_dev_select: ENTER, ata1: device 0, wait 1
ata_dev_identify: do ATA identify
ata_sg_setup_one: mapped buffer of 512 bytes for read
ahci_fill_sg: ENTER

```

```
# cat /proc/modules
```

```

ahci 12181 1 – Loading 0xf8820000
libata 53732 1 ahci, Live 0xf9031000
scsi_mod 131580 2 ahci,libata, Live 0xf905a000
tg3 87808 0 – Live 0xf8860000
uhci_hcd 34856 0 – Live 0xf8828000
usbkbd 7456 0 – Live 0xf881a000
usbcore 121352 3 uhci_hcd,usbkbd, Live 0xf9001000
unix 29260 0 – Live 0xf883e000
fbcon 38496 65 – Live 0xf8833000
font 8352 1 fbcon, Live 0xf8824000
bitblit 5792 1 fbcon, Live 0xf881d000

```

```
# ps -Hel
```

```

F S UID PID PPID C PRI NI ADDR SZ WCHAN TTY TIME CMD
4 S 0 1 0 0 76 0 – 389 wait ? 00:00:00 init
1 S 0 2 1 0 –40 0 – 0 migrat ? 00:00:00 migration/0
1 S 0 3 1 0 94 19 – 0 ksofti ? 00:00:00 ksoftirqd/0
1 S 0 4 1 0 –40 0 – 0 migrat ? 00:00:00 migration/1
1 S 0 5 1 0 94 19 – 0 ksofti ? 00:00:00 ksoftirqd/1
1 S 0 6 1 0 66 –10 – 0 worker ? 00:00:00 events/0
1 S 0 8 6 0 67 –10 – 0 worker ? 00:00:00 khelper
1 S 0 53 6 0 66 –10 – 0 worker ? 00:00:00 kblockd/1
1 S 0 64 6 0 75 0 – 0 pdflus ? 00:00:00 pdflush
1 S 0 67 6 0 66 –10 – 0 worker ? 00:00:00 aio/1
1 S 0 761 6 0 66 –10 – 0 worker ? 00:00:00 ata/0
1 S 0 762 6 0 66 –10 – 0 worker ? 00:00:00 ata/1
1 S 0 7 1 0 65 –10 – 0 worker ? 00:00:00 events/1

```

Linux-Kernel: 2.6.10-rc1 & ahci & IHC6R & 925X & raid1

1 S 0 9 7 0 69 -10 - 0 worker ? 00:00:00 kacpid
1 S 0 52 7 0 66 -10 - 0 worker ? 00:00:00 kblockd/0
1 S 0 63 7 0 80 0 - 0 pdflus ? 00:00:00 pdflush
1 S 0 66 7 0 66 -10 - 0 worker ? 00:00:00 aio/0
1 S 0 65 1 0 77 0 - 0 kswapd ? 00:00:00 kswapd0
1 S 0 653 1 0 81 0 - 0 serio_ ? 00:00:00 kseriod
1 S 0 724 1 0 75 0 - 0 - ? 00:00:00 kirqd
1 S 0 733 1 0 75 0 - 0 hub_th ? 00:00:00 khubd
4 S 0 753 1 0 76 0 - 391 wait ? 00:00:00 sh
4 D 0 760 753 0 78 0 - 379 ata_de ? 00:00:00 modprobe
0 R 0 773 753 0 77 0 - 573 - ? 00:00:00 ps
1 S 0 763 1 0 79 0 - 0 415898 ? 00:00:00 scsi_ah_0
1 S 0 764 1 0 80 0 - 0 415899 ? 00:00:00 scsi_ah_1
1 S 0 765 1 0 81 0 - 0 415907 ? 00:00:00 scsi_ah_2
1 S 0 766 1 0 82 0 - 0 415907 ? 00:00:00 scsi_ah_3

-

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/>