DragonFlyBSD - Bug #732

Live CD problems.

07/22/2007 02:19 AM - luxh

Status:	Closed	**Start date:**	
Priority:	Urgent	**Due date:**	
Assignee:		**% Done:**	0%
Category:		**Estimated time:**	0.00 hour
Target version:			

Description

Hi,

During boot of the live cd this shows up:

Mounting root from cd9660:cd0c
ata1: FAILURE - zero length DMA transfer attempted
acd0: setting up DMA failed
ata1: FAILURE - zero length DMA transfer attempted
acd0: setting up DMA failed
ata1: FAILURE - zero length DMA transfer attempted
acd0: setting up DMA failed
ata1: FAILURE - zero length DMA transfer attempted
acd0: setting up DMA failed
ata1: FAILURE - zero length DMA transfer attempted
acd0: setting up DMA failed
(cd0:ata1:0:0:0): READ(10). CDB: 28 0 0 0 0 10 0 0 0 0
(cd0:ata1:0:0:0): error code 0
(cd0:ata1:0:0:0): cddone: got error 0x5 back
Root mount failed: 5
Mounting root from cd9660:acd0c

Later on, if you start the installer and try to install, or issue
something simple
as 'pwd' this happens:

dscheck(#acd/2): attempt to access non-existant partition
dscheck(#acd/2): attempt to access non-existant partition
vm_fault: pager read error, pid 830 (csh)
pid 830 (csh), uid 0: exited on signal 11 (core dumped)

Here's dmesg.boot:

Copyright (c) 2003, 2004, 2005, 2006, 2007 The DragonFly Project.
Copyright (c) 1992-2003 The FreeBSD Project.
Copyright (c) 1979, 1980, 1983, 1986, 1988, 1989, 1991, 1992, 1993, 1994
The Regents of the University of California. All rights reserved.
DragonFly 1.9.0-DEVELOPMENT #1: Sun Jul 22 02:53:17 CEST 2007
root@leela.:/usr/obj/usr/src/sys/GENERIC
TSC clock: 994726803 Hz, i8254 clock: 1193169 Hz
CPU: Intel Celeron (994.75-MHz 686-class CPU)
Origin = "GenuineIntel" Id = 0x68a Stepping = 10
Features=0x383f9ff<FPU,VME,DE,PSE,TSC,MSR,PAE,MCE,CX8,SEP,MTRR,PGE,MCA,CMOV,PAT,PSE36,MMX,FXSR,SSE>
real memory = 234881024 (229376K bytes)
avail memory = 215498752 (210448K bytes)
Preloaded elf kernel "/kernel" at 0xc07b4000.
Preloaded elf module "/modules/acpi.ko" at 0xc07b4244.
Pentium Pro MTRR support enabled
md0: Malloc disk
pcibios: BIOS version 2.10
Using $PIR table, 8 entries at 0xc00fdf40
ACPI: RSDP @ 0x0xf72b0/0x0014 (v 0 PTLTD)
ACPI: RSDT @ 0x0xdefbfe4/0x0028 (v 1 PTLTD   RSDT   0x06040000 LTP

```
0x00000000)
ACPI: FACP @ 0x0xdefef8c/0x0074 (v  1 VT8606 twisterT 0x06040000 PTL_
0x000F4240)
ACPI: DSDT @ 0x0xdefc00c/0x2F80 (v  1 COMPAL   CY23   0x06040000 MSFT
0x0100000D)
ACPI: FACS @ 0x0xdefffc0/0x0040
npx0: <math processor> on motherboard
npx0: INT 16 interface
Using MMX optimized bcopy/copyin/copyout
acpi0: <PTLTD   RSDT> on motherboard
acpi0: Power Button (fixed)
Warning: ACPI is disabling APM's device.  You can't run both
acpi_timer0: <24-bit timer at 3.579545MHz> port 0x8008-0x800b on acpi0
cpu0: <ACPI CPU (2 Cx states)> on acpi0
acpi_tz0: <Thermal Zone> on acpi0
acpi_lid0: <Control Method Lid Switch> on acpi0
acpi_button0: <Power Button> on acpi0
acpi_ec0: <Embedded Controller: GPE 0x6> port 0x66,0x62 on acpi0
acpi_acad0: <AC Adapter> on acpi0
acpi_cmbat0: <Control Method Battery> on acpi0
legacypci0 on motherboard
pcib0: <Host to PCI bridge> on legacypci0
pci0: <PCI bus> on pcib0
agp0: <VIA 82C694X (Apollo Pro 133A) host to PCI bridge> mem
0xec000000-0xeffffffff at device 0.0 on pci0
pcib1: <PCI to PCI bridge (vendor=1106 device=8605)> at device 1.0 on pci0
pci1: <PCI bus> on pcib1
pci1: <S3 model 8d01 graphics accelerator> at 0.0 irq 10
cbb0: <O2Micro OZ6933 PCI-CardBus Bridge> at device 4.0 on pci0
cardbus0: <CardBus bus> on cbb0
pccard0: <16-bit PCCard bus> on cbb0
pci_cfgintr: 0:4 INTA routed to irq 10
cbb1: <O2Micro OZ6933 PCI-CardBus Bridge> at device 4.1 on pci0
cardbus1: <CardBus bus> on cbb1
pccard1: <16-bit PCCard bus> on cbb1
pci_cfgintr: 0:4 INTB routed to irq 5
isab0: <VIA 82C686 PCI-ISA bridge> at device 7.0 on pci0
isa0: <ISA bus> on isab0
atapci0: <VIA 82C686B UDMA100 controller> port
0x1000-0x100f,0x376,0x170-0x177,0x3f6,0x1f0-0x1f7 at device 7.1 on
pci0
ata0: <ATA channel 0> on atapci0
acpi_cpu: throttling enabled, 16 steps (100% to 6.2%), currently 100.0%
ad0: 19077MB <IC25N020ATDA04 0 DA3OA76A> at ata0-master UDMA100
ata1: <ATA channel 1> on atapci0
acd0: CDRW <TOSHIBA DVD-ROM SD-R2102/1004> at ata1-master UDMA33
uhci0: <VIA 83C572 USB controller> irq 11 at device 7.2 on pci0
usb0: <VIA 83C572 USB controller> on uhci0
usb0: USB revision 1.0
uhub0: 2 ports with 2 removable, self powered
isab1: <PCI to ISA bridge (vendor=1106 device=3057)> at device 7.4 on pci0
pci0: <VIA 82C686 AC97 Audio> at 7.5 irq 5
fwohci0: <Texas Instruments TSB43AB22/A> irq 5 at device 13.0 on pci0
fwohci0: OHCI version 1.10 (ROM=0)
fwohci0: No. of Isochronous channel is 4.
fwohci0: EUI64 00:02:3f:1b:13:00:20:13
fwohci0: Phy 1394a available S400, 2 ports.
fwohci0: Link S400, max_rec 2048 bytes.
firewire0: <IEEE1394(FireWire) bus> on fwohci0
fwe0: <Ethernet over FireWire> on firewire0
fwe0: MAC address: 02:02:3f:00:20:13
sbp0: <SBP-2/SCSI over FireWire> on firewire0
fwohci0: Initiate bus reset
fwohci0: node_id=0xc000ffc0, gen=1, CYCLEMASTER mode
firewire0: 1 nodes, maxhop <= 0, cable IRM = 0 (me)
firewire0: bus manager 0 (me)
pci0: <unknown card> (vendor=0x11c1, dev=0x0450) at 16.0 irq 11
```

dc0: <Accton EN2242 MiniPCI 10/100BaseTX> mem 0xe8004c00-0xe8004fff
irq 11 at device 17.0 on pci0
miibus0: <MII bus> on dc0
ukphy0: <Generic IEEE 802.3u media interface> on miibus0
ukphy0: 10baseT, 10baseT-FDX, 100baseTX, 100baseTX-FDX, auto
dc0: MAC address: 00:90:96:1e:5d:cf
orm0: <Option ROMs> at iomem
0xc0000-0xcbfff,0xcc000-0xcdfff,0xe4000-0xe7fff on isa0
pmtimer0 on isa0
fdc0: <NEC 72065B or clone> at port 0x3f0-0x3f5,0x3f7 irq 6 drq 2 on isa0
fdc0: FIFO enabled, 8 bytes threshold
fd0: <1440-KB 3.5" drive> on fdc0 drive 0
atkbdc0: <Keyboard controller (i8042)> at port 0x60,0x64 on isa0
atkbd0: <AT Keyboard> flags 0x1 irq 1 on atkbdc0
kbd0 at atkbd0
psm0: <PS/2 Mouse> irq 12 on atkbdc0
psm0: model GlidePoint, device ID 0
vga0: <Generic ISA VGA> at port 0x3c0-0x3df iomem 0xa0000-0xbffff on isa0
sc0: <System console> at flags 0x100 on isa0
sc0: VGA <16 virtual consoles, flags=0x300>
sio0: can't drain, serial port might not exist, disabling
sio1: can't drain, serial port might not exist, disabling
ppc0: <Parallel port> at port 0x378-0x37f irq 7 on isa0
ppc0: SMC-like chipset (ECP/EPP/PS2/NIBBLE) in COMPATIBLE mode
ppc0: FIFO with 16/16/8 bytes threshold
ppbus0: <Parallel port bus> on ppc0
plip0: <PLIP network interface> on ppbus0
lpt0: <Printer> on ppbus0
lpt0: Interrupt-driven port
ppi0: <Parallel I/O> on ppbus0
cd0 at ata1 bus 0 target 0 lun 0
cd0: <TOSHIBA DVD-ROM SD-R2102 1004> Removable CD-ROM SCSI-0 device
cd0: 33.000MB/s transfers
cd0: cd present [139184 x 2048 byte records]
Mounting root from cd9660:cd0c
ata1: FAILURE - zero length DMA transfer attempted
acd0: setting up DMA failed
ata1: FAILURE - zero length DMA transfer attempted
acd0: setting up DMA failed
ata1: FAILURE - zero length DMA transfer attempted
acd0: setting up DMA failed
ata1: FAILURE - zero length DMA transfer attempted
acd0: setting up DMA failed
ata1: FAILURE - zero length DMA transfer attempted
acd0: setting up DMA failed
(cd0:ata1:0:0:0): READ(10). CDB: 28 0 0 0 0 10 0 0 0 0
(cd0:ata1:0:0:0): error code 0
(cd0:ata1:0:0:0): cddone: got error 0x5 back
Root mount failed: 5
Mounting root from cd9660:acd0c

Thanks,
Max Herrgaard

History

#1 - 07/22/2007 07:41 PM - dillon

:Hi,
:
:During boot of the live cd this shows up:
:
:Mounting root from cd9660:cd0c
:ata1: FAILURE - zero length DMA transfer attempted
:acd0: setting up DMA failed
:ata1: FAILURE - zero length DMA transfer attempted
:acd0: setting up DMA failed
:ata1: FAILURE - zero length DMA transfer attempted
:acd0: setting up DMA failed

:ata1: FAILURE - zero length DMA transfer attempted
:acd0: setting up DMA failed
:ata1: FAILURE - zero length DMA transfer attempted
:acd0: setting up DMA failed

That is definitely not correct. If you don't mind burning another dead CD could you please try this patch? It will force the system to panic when it tries to do a 0-length I/O. Then get a backtrace from the DDB prompt so we can see the call chain that leads up to the problem.

-Matt

```
Index: ata-dma.c
===================================================================
RCS file: /cvs/src/sys/dev/disk/nata/ata-dma.c,v
retrieving revision 1.4
diff -u -p -r1.4 ata-dma.c
--- ata-dma.c 5 Jun 2007 18:30:40 -0000 1.4
+++ ata-dma.c 22 Jul 2007 19:32:04 -0000
@@ -227,6 +227,7 @@  return EIO;
}
if (!count) {
device_printf(dev, "FAILURE - zero length DMA transfer attempted\n");
+ panic("zero length DMA transfer");
return EIO;
}
if (((uintptr_t)data & (ch->dma->alignment - 1)) ||
Index: atapi-cd.c
===================================================================
RCS file: /cvs/src/sys/dev/disk/nata/atapi-cd.c,v
retrieving revision 1.7
diff -u -p -r1.7 atapi-cd.c
--- atapi-cd.c 3 Jun 2007 04:48:29 -0000 1.7
+++ atapi-cd.c 22 Jul 2007 19:35:03 -0000
@@ -785,6 +785,8 @@  biodone(bp);
return 0;
}

+    KASSERT(bbp->b_bcount != 0, ("acd_strategy: 0-length I/O"));
+
bp->bio_driver_info = cdev;
bbp->b_resid = bbp->b_bcount;

@@ -842,6 +844,8 @@  lba = (bp->bio_offset & 0x00FFFFFFFFFF
}

count = bbp->b_bcount / blocksize;
+    KASSERT(count != 0, ("acd_strategy: 0-length I/O %d bytes vs %d blksize",
+  bbp->b_bcount, blocksize));

if (bbp->b_cmd == BUF_CMD_READ) {
/* if transfer goes beyond range adjust it to be within limits */
```

#2 - 07/23/2007 02:00 AM - tomaz.borstnar

Snapshot from 22th July and basically the same problems:

WinFast K7S741M Series 426XP215 072204
XP2400+

Samsung CRRW+DVD-R

acd0: FAILURE - READ_BIG HARDWARE ERROR asc=0x08 ascq=0x03
vm_fault: pager read error, pid 1 (init)

This string is common on Google so I decided to disable UDMA on secondary channel and got this:

ata1: FAILURE - zero length DMA transfer attempted
acd0: setting up DMA failed
(cd0:ata1:0:0:0): READ(10). CDB: 28 0 0 0 0 10 0 0 0 0
(cd0:ata1:0:0:0): error code 0
(cd0:ata1:0:0:0): cddone: got error 0x5 back
Root mount failed: 5
Mounting root from cd9660:acd0c
cd9660: RockRidge Extension
acd0: FAILURE - READ_BIG HARDWARE ERROR asc=0x08 ascq=0x03
vm_fault: pager read error, pid 1 (swapper)
init died (signal 6, exit 0)
panic: Going nowhere without my init!

With CD-R I got the same results as in the first place (which was done on CD-RW).

I got one backtrace, but did not write it down :(

Tomaž

#3 - 07/23/2007 07:39 AM - luxh

Is this sufficient?

(kgdb) bt
#0 dumpsys () at thread.h:83
#1 0xc02eb11b in boot (howto=260) at /usr/src/sys/kern/kern_shutdown.c:371
#2 0xc02eb8a4 in panic (fmt=0xc0562c05 "from debugger") at
/usr/src/sys/kern/kern_shutdown.c:796
#3 0xc0166a66 in db_panic (addr=-1068477435, have_addr=0, count=-1,
modif=0xc07d6604 "") at /usr/src/sys/ddb/db_command.c:447
#4 0xc01669fb in db_command (last_cmdp=0xc0634690, cmd_table=0x0,
aux_cmd_tablep=0xc05be6f8, aux_cmd_tablep_end=0xc05be714) at

/usr/src/sys/ddb/db_command.c:343
#5 0xc0166adb in db_command_loop () at /usr/src/sys/ddb/db_command.c:469
#6 0xc0169670 in db_trap (type=3, code=0) at /usr/src/sys/ddb/db_trap.c:71
#7 0xc05050a8 in kdb_trap (type=3, code=0, regs=0xc07d6728) at
/usr/src/sys/platform/pc32/i386/db_interface.c:148
#8 0xc0517f20 in trap (frame=0xc07d6728) at
/usr/src/sys/platform/pc32/i386/trap.c:804
#9 0xc05060f6 in calltrap () at /usr/src/sys/platform/pc32/i386/exception.s:783
#10 0xc0505405 in Debugger (msg=0x12 <Address 0x12 out of bounds>) at
cpufunc.h:73
#11 0xc02eb89c in panic (fmt=0xc05652e4 "zero length DMA transfer") at
/usr/src/sys/kern/kern_shutdown.c:794
#12 0xc01a04c8 in ata_dmaload (dev=0xc622bb28, data=0xc0c14000 "",
count=0, dir=2, addr=0x12, entries=0x0) at
/usr/src/sys/dev/disk/nata/ata-dma.c:230
#13 0xc01a0ca0 in ata_begin_transaction (request=0xc0b2b190) at
/usr/src/sys/dev/disk/nata/ata-lowlevel.c:180
#14 0xc01a4710 in ata_start (dev=0xc622bb28) at
/usr/src/sys/dev/disk/nata/ata-queue.c:214
#15 0xc01a4376 in ata_queue_request (request=0xc0b2b190) at
/usr/src/sys/dev/disk/nata/ata-queue.c:95
#16 0xc01b386d in atapi_action (sim=0xc0b03ba8, ccb=0x0) at
/usr/src/sys/dev/disk/nata/atapi-cam.c:602
#17 0xc014be68 in xpt_run_dev_sendq (bus=0xc0af2c90) at
/usr/src/sys/bus/cam/cam_xpt.c:3640
#18 0xc014b111 in xpt_action (start_ccb=0xca67a140) at
/usr/src/sys/bus/cam/cam_xpt.c:2809
#19 0xc015c90b in cdstart (periph=0xc0b03bf8, start_ccb=0xca67a140) at
/usr/src/sys/bus/cam/scsi/scsi_cd.c:1584
#20 0xc014bbfc in xpt_run_dev_allocq (bus=0xc0af2c90) at
/usr/src/sys/bus/cam/cam_xpt.c:3507
#21 0xc014baa1 in xpt_schedule (perph=0xc0b03bf8,
new_priority=3232705680) at /usr/src/sys/bus/cam/cam_xpt.c:3402
#22 0xc015c5d2 in cdstrategy (ap=0x12) at
/usr/src/sys/bus/cam/scsi/scsi_cd.c:1496
#23 0xc02d456d in dev_dstrategy (dev=0xc0b25fe8, bio=0x20) at
/usr/src/sys/kern/kern_device.c:238
#24 0xc030094e in diskstrategy (ap=0x12) at /usr/src/sys/kern/subr_disk.c:423
#25 0xc02d456d in dev_dstrategy (dev=0xc0b26f18, bio=0x20) at
/usr/src/sys/kern/kern_device.c:238
#26 0xc034824d in spec_strategy (ap=0xc65f5d18) at
/usr/src/sys/vfs/specfs/spec_vnops.c:544
#27 0xc034328b in vop_strategy (ops=0x0, vp=0xc0b6b188,
bio=0xc0ba3e30) at /usr/src/sys/kern/vfs_vopops.c:659
#28 0xc0329357 in vn_strategy (vp=0xc0b6b188, bio=0x20) at
/usr/src/sys/kern/vfs_bio.c:2779
#29 0xc0326140 in bread (vp=0xc0b6b188, loffset=) at
/usr/src/sys/kern/vfs_bio.c:630
#30 0xc02c2f4e in iso_mountfs (devvp=0xc0b6b188, mp=0xc0b84020,
argp=0xc07d6b64) at /usr/src/sys/vfs/isofs/cd9660/cd9660_vfsops.c:321
#31 0xc02c2ad8 in iso_mountroot (mp=0xc0b84020) at
/usr/src/sys/vfs/isofs/cd9660/cd9660_vfsops.c:165
#32 0xc02c2b36 in cd9660_mount (mp=0xc0b84020, path=0x0, data=0x12
<Address 0x12 out of bounds>, cred=0xc0b16120) at
/usr/src/sys/vfs/isofs/cd9660/cd9660_vfsops.c:189
#33 0xc032f650 in vfs_mountroot_try (mountfrom=0xc0577e53
"cd9660:cd0c") at /usr/src/sys/kern/vfs_conf.c:244
#34 0xc032f3f4 in vfs_mountroot (junk=0x0) at /usr/src/sys/kern/vfs_conf.c:143
#35 0xc02c7cd1 in mi_startup () at /usr/src/sys/kern/init_main.c:234
#36 0xc014098b in begin () at /usr/src/sys/platform/pc32/i386/locore.s:337

#4 - 07/23/2007 05:21 PM - dillon

:Is this sufficient?

This helps. I'm researching it now. Somehow a 2K read got turned
into a 0 block read.

-Matt

#5 - 07/23/2007 08:11 PM - dillon

Ok, all the live CD problems should now theoretically be fixed. I'm
doing a full nrelease build to test it.

I found numerous additional bugs while tracking down this one,
including problems tracking the open count on the device. This
typically prevented the CD from being ejected. /dev/cd0 also had the
wrong minor number, and the generic disk layer also failed to properly
track opens when an open failed (which I found when I couldn't
eject the CD after trying to open the incorrect /dev/cd0 device).

The main issue was that the CAM layer failed to set the si_iosize_max
field in the device structure. This caused the strategy code to
try to break I/O's down into 0-sized chunks, resulting in an endless
loop and/or panic.

NATA's CAM interface also failed to properly set transfer length
limits based on the DMA transfer limit and instead just used 65535
unconditionally. This resulted in accesses via the SCSI device
(/dev/cd0) potentially making I/O requests that were too large.

-Matt

#6 - 07/24/2007 12:12 PM - luxh

...
> -Matt

I did one too and it seems to work. Thanks!

Max

#7 - 07/24/2007 12:16 PM - luxh

Commit msgs:
http://leaf.dragonflybsd.org/mailarchive/commits/2007-07/msg00198.html
http://leaf.dragonflybsd.org/mailarchive/commits/2007-07/msg00200.html