

Re: /dev/null and KSE panic 100% reproducible

Source: <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/current/2003-05/0331.html>

From: Julian Elischer (julian_at_elischer.org)

Date: 05/20/03

Date: Tue, 20 May 2003 11:39:11 -0700 (PDT)

To: "Daniel C. Sobral" <dcsoip.com.br>

On Tue, 20 May 2003, Daniel C. Sobral wrote:

> *Or so it seems. If I do a make install in*
> */usr/ports/emulators/linux_base, panic happens. Alas, since my first*
> *panic yesterday was KSE and was during portupgrade and I have no*
> *linux_base presently installed, I suspect this is what caused that first*
> *panic.*
>
> *Julian, since I now have a _reproducible_ KSE panic... what do you want*
> *me to do? :-)*
>
> *New backtrace, for your delight and enjoyment.*
>
> *GNU gdb 5.2.1 (FreeBSD)*
> *Copyright 2002 Free Software Foundation, Inc.*
> *GDB is free software, covered by the GNU General Public License, and you are*
> *welcome to change it and/or distribute copies of it under certain*
> *conditions.*
> *Type "show copying" to see the conditions.*
> *There is absolutely no warranty for GDB. Type "show warranty" for details.*
> *This GDB was configured as "i386-undermydesk-freebsd"...*
> *panic: KSE not on run queue*
> *panic messages:*
> *----*
> *panic: No strategy on dev null responsible for buffer 0xc78084f8*

this is odd.

it is very hard to work out which panic is occurring first..
is it the dev null or the KSE panic?

can you get a serial console so we can be sure about this?

the stack trace you showed is in the context of a clock interrupt,
(or at least, SOME interrupt). The possibility is that the clock
interrupt is recalculating priorities but somehow it's happening when

freebsd-current: Re: /dev/null and KSE panic 100% reproducible

the system is already messing up it's scheduling data..

Someone put the following code into kern/kern_switch.c

```
/*
 * Only allow non system threads to run in panic
 * if they are the one we are tracing. (I think.. [JRE])
 */
if (panicstr && ((td->td_proc->p_flag & P_SYSTEM) == 0 &&
    (td->td_flags & TDF_INPANIC) == 0))
    goto retry;
```

It has the effect of throwing away threads that it has taken off the run queue if we are in a panic.

at a later time anything that goes through these threads will assume they are still on the run queue and panic because they are not..

try the following:

change it to:

```
if (panicstr && ((td->td_proc->p_flag & P_SYSTEM) == 0 &&
    (td->td_flags & TDF_INPANIC) == 0)) {
    /* note that it is no longer on the run queue */
    TD_SET_CAN_RUN(td);
    goto retry;
}
```

if it fails you may try TD_SET_SUSPENDED(td) instead, but I think this is better.

```
>
>
> syncing disks, buffers remaining... 2230 2230 2230 2229 2229 2229 2229
> 2229 panic: KSE not on run queue
> Uptime: 50m29s
> Dumping 255 MB
> ata0: resetting devices ..
> done
> 16 32 48 64 80 96 112 128 144 160 176[CTRL-C to abort] [CTRL-C to
> abort] [CTRL-C to abort] [CTRL-C to abort] [CTRL-C to abort] [CTRL-C to
> abort] 192 208 224 240Copyright (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.
> FreeBSD 5.1-BETA #30: Mon May 19 21:40:33 BRT 2003
> root@dcs:/usr/obj/usr/src/sys/DCS
> Preloaded elf kernel "/boot/kernel/kernel" at 0xc04c0000.
> Preloaded elf module "/boot/kernel/snd_cmi.ko" at 0xc04c0228.
> Preloaded elf module "/boot/kernel/snd_pcm.ko" at 0xc04c02d4.
> Preloaded elf module "/boot/kernel/mac_biba.ko" at 0xc04c0380.
> Preloaded elf module "/boot/kernel/mac_mls.ko" at 0xc04c0430.
```

Re: /dev/null and KSE panic 100% reproducible

freebsd-current: Re: /dev/null and KSE panic 100% reproducible

```
> Preloaded elf module "/boot/kernel/acpi.ko" at 0xc04c04dc.
> Timecounter "i8254" frequency 1193182 Hz
> Timecounter "TSC" frequency 1007051981 Hz
> CPU: Intel(R) Celeron(TM) CPU 1000MHz (1007.05-MHz
> 686-class CPU)
> Origin = "GenuineIntel" Id = 0x6b1 Stepping = 1
>
>
Features=0x383f9ff<FPU,VME,DE,PSE,TSC,MSR,PAE,MCE,CX8,SEP,MTRR,PGE,MCA,CMOV,PAT,PSE36,MMX,
> real memory = 268353536 (255 MB)
> avail memory = 255438848 (243 MB)
> Security policy loaded: TrustedBSD MAC/Biba (mac_biba)
> Security policy loaded: TrustedBSD MAC/MLS (mac_mls)
> Pentium Pro MTRR support enabled
> VESA: v3.0, 4096k memory, flags:0x1, mode table:0xc03ca722 (1000022)
> VESA: NVidia
> npx0: <math processor> on motherboard
> npx0: INT 16 interface
> acpi0: <ASUS TUV4X > on motherboard
> pcibios: BIOS version 2.10
> Using $PIR table, 9 entries at 0xc00f12d0
> acpi0: power button is handled as a fixed feature programming model.
> Timecounter "ACPI-fast" frequency 3579545 Hz
> acpi_timer0: <24-bit timer at 3.579545MHz> port 0xe408-0xe40b on acpi0
> acpi_cpu0: <CPU> on acpi0
> acpi_button0: <Power Button> on acpi0
> pcib0: <ACPI Host-PCI bridge> port 0xcf8-0xcff on acpi0
> pci0: <ACPI PCI bus> on pcib0
> agp0: <VIA 82C691 (Apollo Pro) host to PCI bridge> mem
> 0xf8000000-0xfbffffff at device 0.0 on pci0
> pcib1: <ACPI PCI-PCI bridge> at device 1.0 on pci0
> pcib1: could not get PCI interrupt routing table for \_SB_.PCI0.AGP_ -
> AE_NOT_FOUND
> pci1: <ACPI PCI bus> on pcib1
> pci1: <display, VGA> at device 0.0 (no driver attached)
> isab0: <PCI-ISA bridge> at device 4.0 on pci0
> isa0: <ISA bus> on isab0
> atapci0: <VIA 82C686B UDMA100 controller> port 0xd800-0xd80f at device
> 4.1 on pci0
> ata0: at 0x1f0 irq 14 on atapci0
> ata1: at 0x170 irq 15 on atapci0
> uhci0: <VIA 83C572 USB controller> port 0xd400-0xd41f irq 5 at device
> 4.2 on pci0
> usb0: <VIA 83C572 USB controller> on uhci0
> usb0: USB revision 1.0
> uhub0: VIA UHCI root hub, class 9/0, rev 1.00/1.00, addr 1
> uhub0: 2 ports with 2 removable, self powered
> uhub0: port error, restarting port 1
> uhub0: port error, giving up port 1
> ugen0: AKS eToken R2 2242, rev 1.00/1.00, addr 2
> uhub0: port error, restarting port 2
```

freebsd-current: Re: /dev/null and KSE panic 100% reproducible

```
> uhub0: port error, giving up port 2
> uhci1: <VIA 83C572 USB controller> port 0xd000-0xd01f irq 5 at device
> 4.3 on pci0
> usb1: <VIA 83C572 USB controller> on uhci1
> usb1: USB revision 1.0
> uhub1: VIA UHCI root hub, class 9/0, rev 1.00/1.00, addr 1
> uhub1: 2 ports with 2 removable, self powered
> uhub1: port error, restarting port 1
> uhub1: port error, giving up port 1
> uhub1: port error, restarting port 2
> uhub1: port error, giving up port 2
> pcm0: <CMedia CMI8738> port 0xb800-0xb8ff at device 5.0 on pci0
> pcib0: slot 5 INTA is routed to irq 5
> fxp0: <Intel 82557/8/9 EtherExpress Pro/100(B) Ethernet> port
> 0xb400-0xb43f mem 0xf3000000-0xf30fffff,0xf3800000-0xf3800fff irq 10 at
> device 9.0 on pci0
> fxp0: Ethernet address 00:02:b3:ae:0d:ea
> miibus0: <MII bus> on fxp0
> inphy0: <i82555 10/100 media interface> on miibus0
> inphy0: 10baseT, 10baseT-FDX, 100baseTX, 100baseTX-FDX, auto
> fdc0: <Enhanced floppy controller (i82077, NE72065 or clone)> port
> 0x3f7,0x3f2-0x3f5 irq 6 drq 2 on acpi0
> fdc0: FIFO enabled, 8 bytes threshold
> fd0: <1440-KB 3.5" drive> on fdc0 drive 0
> ppc0 port 0x378-0x37f irq 7 on acpi0
> ppc0: Generic chipset (EPP/NIBBLE) in COMPATIBLE mode
> ppbus0: <Parallel port bus> on ppc0
> lpt0: <Printer> on ppbus0
> lpt0: Interrupt-driven port
> ppi0: <Parallel I/O> on ppbus0
> sio0 port 0x3f8-0x3ff irq 4 on acpi0
> sio0: type 16550A
> sio1 port 0x2f8-0x2ff irq 3 on acpi0
> sio1: type 16550A
> atkbd0: <Keyboard controller (i8042)> port 0x64,0x60 irq 1 on acpi0
> atkbd0: <AT Keyboard> flags 0x1 irq 1 on atkbd0
> kbd0 at atkbd0
> psm0: <PS/2 Mouse> irq 12 on atkbd0
> psm0: model Generic PS/2 mouse, device ID 0
> orm0: <Option ROMs> at iomem 0xd0000-0xd0fff,0xc0000-0xcffff on isa0
> pmtimer0 on isa0
> sc0: <System console> at flags 0x100 on isa0
> sc0: VGA <16 virtual consoles, flags=0x300>
> vga0: <Generic ISA VGA> at port 0x3c0-0x3df iomem 0xa0000-0xbffff on isa0
> Timecounters tick every 1.000 msec
> acpi_cpu: throttling enabled, 16 steps (100% to 6.2%), currently 100.0%
> ad0: 19092MB <ST320410A> [38792/16/63] at ata0-master UDMA100
> acd0: CDROM <CRD-8481B> at ata1-master PIO4
> Mounting root from ufs:/dev/ad0s2a
> WARNING: / was not properly dismounted
> lock order reversal
```

Re: /dev/null and KSE panic 100% reproducible

freebsd-current: Re: /dev/null and KSE panic 100% reproducible

```
> 1st 0xc25fb840 pcm0 (sound softc) @ /usr/src/sys/dev/sound/pci/cmi.c:520
> 2nd 0xc25fb440 pcm0:play:0 (pcm channel) @
> /usr/src/sys/dev/sound/pcm/channel.c:440
> Stack backtrace:
> panic: No strategy on dev null responsible for buffer 0xc776ab90
>
>
> syncing disks, buffers remaining... 2230 2230 2229 2229 2229 2228 2228
> 2228 2228 2228 2228 panic: KSE not on run queue
> Uptime: 26m44s
> Dumping 255 MB
> ata0: resetting devices ..
> done
> 16 32 48 64 80 96 112 128 144 160 176 192 208 224 240
> ---
> Reading symbols from /boot/kernel/snd_cmi.ko...done.
> Loaded symbols for /boot/kernel/snd_cmi.ko
> Reading symbols from /boot/kernel/snd_pcm.ko...done.
> Loaded symbols for /boot/kernel/snd_pcm.ko
> Reading symbols from
> /usr/obj/usr/src/sys/DCS/modules/usr/src/sys/modules/mac_biba/mac_biba.ko.debug...done.
> Loaded symbols for
> /usr/obj/usr/src/sys/DCS/modules/usr/src/sys/modules/mac_biba/mac_biba.ko.debug
> Reading symbols from
> /usr/obj/usr/src/sys/DCS/modules/usr/src/sys/modules/mac_mls/mac_mls.ko.debug...done.
> Loaded symbols for
> /usr/obj/usr/src/sys/DCS/modules/usr/src/sys/modules/mac_mls/mac_mls.ko.debug
> Reading symbols from
> /usr/obj/usr/src/sys/DCS/modules/usr/src/sys/modules/acpi/acpi.ko.debug...done.
> Loaded symbols for
> /usr/obj/usr/src/sys/DCS/modules/usr/src/sys/modules/acpi/acpi.ko.debug
> Reading symbols from /boot/kernel/green_saver.ko...done.
> Loaded symbols for /boot/kernel/green_saver.ko
> Reading symbols from
> /usr/obj/usr/src/sys/DCS/modules/usr/src/sys/modules/linux/linux.ko.debug...done.
> Loaded symbols for
> /usr/obj/usr/src/sys/DCS/modules/usr/src/sys/modules/linux/linux.ko.debug
> #0 doadump () at /usr/src/sys/kern/kern_shutdown.c:238
> 238 dumping++;
> (kgdb) bt full
> #0 doadump () at /usr/src/sys/kern/kern_shutdown.c:238
> No locals.
> #1 0xc01e7353 in boot (howto=260) at /usr/src/sys/kern/kern_shutdown.c:370
> No locals.
> #2 0xc01e769b in panic () at /usr/src/sys/kern/kern_shutdown.c:543
> td = (struct thread *) 0xc0ecbab0
> bootopt = 260
> newpanic = 0
> buf = "KSE not on run queue\0ll responsible for buffer
> 0xc776ab90\n", '\0' <repeats 197 times>
> #3 0xc01fa832 in sched_rem (ke=0xc0ecd4c0) at
```

Re: /dev/null and KSE panic 100% reproducible

freebsd-current: Re: /dev/null and KSE panic 100% reproducible

```
> /usr/src/sys/kern/sched_4bsd.c:639
> No locals.
> #4 0xc01ecfd3 in adjustrunqueue (td=0xc0ecd4c0, newpri=180) at
> /usr/src/sys/kern/kern_switch.c:295
> kg = (struct ksegrp *) 0xc0ed0b00
> ke = (struct kse *) 0xc0ecd4c0
> #5 0xc01fa580 in sched_prio (td=0xc0ed24c0, prio=180 '') at
> /usr/src/sys/kern/sched_4bsd.c:545
> No locals.
> #6 0xc01f9fe9 in schedcpu (arg=0x0) at /usr/src/sys/kern/sched_4bsd.c:337
> loadfac = 1944
> td = (struct thread *) 0xc0ed24c0
> p = (struct proc *) 0xc26425a0
> ke = (struct kse *) 0x0
> kg = (struct ksegrp *) 0xc0ed0b00
> realstathz = 128
> awake = 1
> #7 0xc01f69ec in softclock (dummy=0x0) at
> /usr/src/sys/kern/kern_timeout.c:195
> c_func = (void (*)(void *)) 0xc01f9e30 <schedcpu>
> c_arg = (void *) 0x0
> c_flags = 14
> c = (struct callout *) 0x0
> bucket = (struct callout_tailq *) 0xc771d9a0
> curticks = 1604096
> steps = 14
> #8 0xc01c59c2 in ithread_loop (arg=0xc0ec8f00) at
> /usr/src/sys/kern/kern_intr.c:537
> ithd = (struct ithd *) 0xc0ec8f00
> ih = (struct intrhand *) 0xc0ec0440
> td = (struct thread *) 0xc0ecbab0
> p = (struct proc *) 0xc0eca780
> ---Type <return> to continue, or q <return> to quit---
> #9 0xc01c49b0 in fork_exit (callout=0xc0ec0440, arg=0x0, frame=0x0) at
> /usr/src/sys/kern/kern_fork.c:768
> td = (struct thread *) 0x0
> p = (struct proc *) 0xc0ec8f00
>
>
> --
> Daniel C. Sobral (8-DCS)
> Gerencia de Operacoes
> Divisao de Comunicacao de Dados
> Coordenacao de Seguranca
> VIVO Centro Oeste Norte
> Fones: 55-61-313-7654/Cel: 55-61-9618-0904
> E-mail: Daniel.Capo@tco.net.br
> Daniel.Sobral@tcoip.com.br
> dcs@tcoip.com.br
>
> Outros:
```

Re: /dev/null and KSE panic 100% reproducible

freebsd-current: Re: /dev/null and KSE panic 100% reproducible

> *dcs@newsguy.com*
> *dcs@freebsd.org*
> *capo@notorious.bsdcconspiracy.net*
>
> *God is Dead*
> -- *Nietzsche*
> *Nietzsche is Dead*
> -- *God*
> *Nietzsche is God*
> -- *The Dead*
>
>

freebsd-current@freebsd.org mailing list

<http://lists.freebsd.org/mailman/listinfo/freebsd-current>

To unsubscribe, send any mail to "freebsd-current-unsubscribe@freebsd.org"