

Re: Odd performance problems after upgrade from 4.11 to 6.0–Stable

Source: <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/stable/2005-12/msg00369.html>

- *From:* "Kevin Oberman" <oberman@xxxxxx>
 - *Date:* Wed, 14 Dec 2005 16:26:18 -0800
-

> Date: Wed, 14 Dec 2005 16:17:38 -0700
> From: Scott Long <scottl@xxxxxxxxxx>
>
> Kevin Oberman wrote:
>
>> I recently upgraded my last desktop system from 4.11–Stable to
>> 6.0–Stable. I did an update to 5.3 then to RELENG_5,
>> RELENG_6_0_0_RELEASE and on to RELENG6_0.
>>
>> This system has been updated regularly from the days of at least 4.1.
>> The hardware is a 1GHz PIII with an ICH2 chipset. 256 MB of memory.
>>
>> After the upgrade, the system performs poorly. Specifically, it seems to
>> be blocking on I/O. If I have a dump or bsdtar running, a CPU intensive
>> job produced minimal CPU utilization and little progress. The system is
>> 70% or more idle all of the time. X is particularly sluggish. Many times
>> the system will simply appear to be locked up for seconds at a time like
>> all of the I/O is synchronous and the entire system is blocking on GIANT
>> until the disk write finishes. A level 0 dump of the system drive piped
>> through gzip and over the network took about 10 hours. Now it is 17
>> hours. My network link only ran at 3.4 Mbps (yes, that's bits, not
>> Bytes) over a 100 Mbps connection to the storage server. The rate was
>> much higher under 4.11.
>>
>> My attempts to track down the source of the problem simply failed. I
>> have examined my kernel configuration and found nothing out of the
>> ordinary. In fact, the the configuration varies from GENERIC only in
>> devices (both added and deleted), dropping 486 and 586 CPU code and a
>> larger SC_HISTORY. I have nothing in sysctl.conf and only enable DMA on
>> my ATAPI devices in loader.conf. I have confirmed that the disks are
>> running ATA100.
>>
>> Any suggestions on where to look? At this point I am baffled.
>
> The bufdaemon is probably causing some of the Giant contention.
> However, what it is contending with is the real question. Is
> your network driver and stack running MPSAFE? At least a 'dmesg'
> dump is needed here.

Re: Odd performance problems after upgrade from 4.11 to 6.0–Stable

>
> Also, taking out CPU_I586 is usually a bad idea. It offers no
> performance penalties (unlike CPU_I386 and maybe CPU_I486), but
> enables things like optimized bcopy.

Ahh, This is the sort of thing I never realized. Is there anything in the handbook that covers this? I had always been under the impression that CPU_I686 enabled all things that the 686 was capable of. I will build a new kernel to add that back in.

I am attaching a dmesg. I do have a few of drivers (uhci, pcm, psm, atkbd0 and ichsmb) that are still marked as GIANT–LOCKED, but I'm not using the USB very often. And I'm not using pcm or ichsmb during the dump, either. I think everyone has the mouse and keyboard under GIANT, but I can't really see those as a problem, either.

I will attach both my kernel configuration and my dmesg in hopes that I'm missing something obvious.

Thanks!

—

R. Kevin Oberman, Network Engineer
Energy Sciences Network (ESnet)
Ernest O. Lawrence Berkeley National Laboratory (Berkeley Lab)
E–mail: oberman@xxxxxx Phone: +1 510 486–8634

```
## PTAVV -- Modified ESnet kernel configuration file for FreeBSD/i386
#         running on TWC Dell PIII system.
#
# For more information on this file, please read the handbook section on
# Kernel Configuration Files:
#
# http://www.FreeBSD.org/doc/en\_US.ISO8859-1/books/handbook/kernelconfig-config.html
#
# The handbook is also available locally in /usr/share/doc/handbook
# if you've installed the doc distribution, otherwise always see the
# FreeBSD World Wide Web server (http://www.FreeBSD.org/) for the
# latest information.
#
# An exhaustive list of options and more detailed explanations of the
# device lines is also present in the ../../conf/NOTES and NOTES files.
# If you are in doubt as to the purpose or necessity of a line, check first
# in NOTES.
#
# $FreeBSD: src/sys/i386/conf/GENERIC,v 1.429.2.4 2005/10/28 19:21:27 jhb Exp $

machine      i386
cpu          I486_CPU
cpu          I586_CPU
cpu          I686_CPU
ident       PTAVV

# To statically compile in device wiring instead of /boot/device.hints
#hints      "GENERIC.hints"          # Default places to look for devices.
```

Re: Odd performance problems after upgrade from 4.11 to 6.0–Stable

```
makeoptions      DEBUG=-g                # Build kernel with gdb(1) debug symbols

#options         SCHED_ULE                # ULE scheduler
options          SCHED_4BSD            # 4BSD scheduler
options          PREEMPTION            # Enable kernel thread preemption
options          INET                  # InterNETworking
options          INET6                  # IPv6 communications protocols
options          FFS                    # Berkeley Fast Filesystem
options          SOFTUPDATES            # Enable FFS soft updates support
options          UFS_ACL                # Support for access control lists
options          UFS_DIRHASH           # Improve performance on big directories
#options         MD_ROOT                # MD is a potential root device
#options         NFSCLIENT             # Network Filesystem Client
#options         NFSSERVER             # Network Filesystem Server
#options         NFS_ROOT               # NFS usable as /, requires NFSCLIENT
options          MSDOSFS                # MSDOS Filesystem
options          CD9660                 # ISO 9660 Filesystem
options          PROCFS                 # Process filesystem (requires PSEUDOFS)
options          PSEUDOFS               # Pseudo-filesystem framework
options          GEOM_GPT                # GUID Partition Tables.
options          COMPAT_43              # Compatible with BSD 4.3 [KEEP THIS!]
options          COMPAT_FREEBSD4        # Compatible with FreeBSD4
options          COMPAT_FREEBSD5        # Compatible with FreeBSD5
options          SCSI_DELAY=5000        # Delay (in ms) before probing SCSI
options          KTRACE                 # ktrace(1) support
options          SYSVSHM                # SYSV-style shared memory
options          SYSVMSG                # SYSV-style message queues
options          SYSVSEM                # SYSV-style semaphores
options          _KPOSIX_PRIORITY_SCHEDULING # POSIX P1003_1B real-time extensions
options          KBD_INSTALL_CDEV       # install a CDEV entry in /dev
#options         AHC_REG_PRETTY_PRINT   # Print register bitfields in debug
#options         AHD_REG_PRETTY_PRINT   # Print register bitfields in debug
options          ADAPTIVE_GIANT         # Giant mutex is adaptive.

#device          apic                  # I/O APIC

# Bus support.
#device          eisa
device           pci

# Floppy drives
device           fd

# ATA and ATAPI devices
device           ata
device           atadisk                # ATA disk drives
#device          ataraid                # ATA RAID drives
device           atapicd                # ATAPI CDROM drives
#device          atapifd                # ATAPI floppy drives
#device          atapist                # ATAPI tape drives
device           atapicam                # emulate ATAPI devices as SCSI ditto via CAM
# needs CAM to be present (scbus & pass)
options          ATA_STATIC_ID          # Static device numbering

# SCSI peripherals
device           scbus                  # SCSI bus (required for SCSI)
#device          ch                      # SCSI media changers
device           da                      # Direct Access (disks)
#device          sa                      # Sequential Access (tape etc)
```

Re: Odd performance problems after upgrade from 4.11 to 6.0–Stable

```
device      cd          # CD
device      pass        # Passthrough device (direct SCSI access)
device      ses         # SCSI Environmental Services (and SAF-TE)

# atkbd0 controls both the keyboard and the PS/2 mouse
device      atkbd       # AT keyboard controller
device      atkbd       # AT keyboard
device      psm         # PS/2 mouse

device      vga         # VGA video card driver

device      splash      # Splash screen and screen saver support

# syscons is the default console driver, resembling an SCO console
device      sc
options     SC_HISTORY_SIZE=2000

# Enable this for the pcvt (VT220 compatible) console driver
#device     vt
#options    XSERVER      # support for X server on a vt console
#options    FAT_CURSOR   # start with block cursor

device      agp         # support several AGP chipsets

# Power management support (see NOTES for more options)
#device     apm
# Add suspend/resume support for the i8254.
device      pmtimer

# Serial (COM) ports
device      sio         # 8250, 16[45]50 based serial ports

# Parallel port
#device     ppc
#device     ppbus       # Parallel port bus (required)
#device     lpt         # Printer
#device     plip        # TCP/IP over parallel
#device     ppi         # Parallel port interface device
#device     vpo         # Requires scbus and da

# If you've got a "dumb" serial or parallel PCI card that is
# supported by the puc(4) glue driver, uncomment the following
# line to enable it (connects to the sio and/or ppc drivers):
#device     puc

# PCI Ethernet NICs that use the common MII bus controller code.
# NOTE: Be sure to keep the 'device miibus' line in order to use these NICs!
device      miibus     # MII bus support
device      xl         # 3Com 3c90x (``Boomerang'', ``Cyclone'')

# Pseudo devices.
device      loop       # Network loopback
device      random     # Entropy device
device      ether      # Ethernet support
#device     sl         # Kernel SLIP
#device     ppp        # Kernel PPP
device      tun        # Packet tunnel.
device      pty        # Pseudo-ttys (telnet etc)
device      md         # Memory "disks"
device      gif        # IPv6 and IPv4 tunneling
device      faith      # IPv6-to-IPv4 relaying (translation)
```

Re: Odd performance problems after upgrade from 4.11 to 6.0–Stable

```
# The `bpf' device enables the Berkeley Packet Filter.
# Be aware of the administrative consequences of enabling this!
# Note that 'bpf' is required for DHCP.
device      bpf          # Berkeley packet filter

# USB support
device      uhci         # UHCI PCI->USB interface
device      ohci         # OHCI PCI->USB interface
device      ehci         # EHCI PCI->USB interface (USB 2.0)
device      usb          # USB Bus (required)
#device     udbp         # USB Double Bulk Pipe devices
device      ugen         # Generic
device      uhid         # "Human Interface Devices"
device      ukbd         # Keyboard
device      ulpt         # Printer
device      umass        # Disks/Mass storage - Requires scbus and da
device      ums          # Mouse
#device     ural         # Ralink Technology RT2500USB wireless NICs
#device     urio         # Diamond Rio 500 MP3 player
device      uscanner     # Scanners
# USB Ethernet, requires miibus
#device     aue          # ADMtek USB Ethernet
#device     axe          # ASIX Electronics USB Ethernet
#device     cdce         # Generic USB over Ethernet
#device     cue          # CATC USB Ethernet
#device     kue          # Kawasaki LSI USB Ethernet
#device     rue          # RealTek RTL8150 USB Ethernet

# FireWire support
#device     firewire     # FireWire bus code
#device     sbp          # SCSI over FireWire (Requires scbus and da)
#device     fwe          # Ethernet over FireWire (non-standard!)

# Sound card
device      smbus        #
device      ichsmb       #
device      smb          #
device      sound        #
device      snd_ich      #

# Video capture card
device      iicbus       #
device      iicbb        #
device      bktr         #
options    BROOKTREE_SYSTEM_DEFAULT=BROOKTREE_NTSC
```

Copyright (c) 1992-2005 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 6.0-STABLE #0: Sun Dec 11 23:18:16 PST 2005

oberman@xxxxxxxxxxxxxx:/usr/obj/usr/src/sys/PTAVV

Timecounter "i8254" frequency 1193182 Hz quality 0

CPU: Intel Pentium III (996.77-MHz 686-class CPU)

Origin = "GenuineIntel" Id = 0x68a Stepping = 10

Features=0x383fbff<FPU,VME,DE,PSE,TSC,MSR,PAE,MCE,CX8,APIC,SEP,MTRR,PGE,MCA,CMOV,PAT,PSE36,MMX,

real memory = 267350016 (254 MB)

avail memory = 256102400 (244 MB)

npx0: [FAST]

npx0: <math processor> on motherboard

npx0: INT 16 interface

acpi0: <DELL GX150 > on motherboard

Re: Odd performance problems after upgrade from 4.11 to 6.0–Stable

```
acpi0: Power Button (fixed)
Timecounter "ACPI-fast" frequency 3579545 Hz quality 1000
acpi_timer0: <24-bit timer at 3.579545MHz> port 0x808-0x80b on acpi0
cpu0: <ACPI 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
agg0: <Intel 82815 (i815 GMCH) host to PCI bridge> mem 0xf4000000-0xf7ffffff at device 0.0 on pci0
pcib1: <PCI-PCI bridge> at device 1.0 on pci0
pci1: <PCI bus> on pcib1
pci1: <display, VGA> at device 0.0 (no driver attached)
pcib2: <ACPI PCI-PCI bridge> at device 30.0 on pci0
pci2: <ACPI PCI bus> on pcib2
xl0: <3Com 3c905C-TX Fast Etherlink XL> port 0xec80-0xecff mem 0xfbfff000-0xfbfff07f irq 11 at device 0.0 on pci0
miibus0: <MII bus> on xl0
ukphy0: <Generic IEEE 802.3u media interface> on miibus0
ukphy0: 10baseT, 10baseT-FDX, 100baseTX, 100baseTX-FDX, auto
xl0: Ethernet address: 00:06:5b:80:b1:fc
isab0: <PCI-ISA bridge> at device 31.0 on pci0
isa0: <ISA bus> on isab0
atapci0: <Intel ICH2 UDMA100 controller> port 0x1f0-0x1f7,0x3f6,0x170-0x177,0x376,0xffa0-0xffaf on pci0
ata0: <ATA channel 0> on atapci0
ata1: <ATA channel 1> on atapci0
uhci0: <Intel 82801BA/BAM (ICH2) USB controller USB-A> port 0xff80-0xff9f irq 11 at device 31.2 on pci0
uhci0: [GIANT-LOCKED]
usb0: <Intel 82801BA/BAM (ICH2) USB controller USB-A> on uhci0
usb0: USB revision 1.0
uhub0: Intel UHCI root hub, class 9/0, rev 1.00/1.00, addr 1
uhub0: 2 ports with 2 removable, self powered
ichsmb0: <Intel 82801BA (ICH2) SMBus controller> port 0xdcd0-0xdcdf irq 10 at device 31.3 on pci0
ichsmb0: [GIANT-LOCKED]
smbus0: <System Management Bus> on ichsmb0
smb0: <SMBus generic I/O> on smbus0
uhci1: <Intel 82801BA/BAM (ICH2) USB controller USB-B> port 0xff60-0xff7f irq 9 at device 31.4 on pci0
uhci1: [GIANT-LOCKED]
usb1: <Intel 82801BA/BAM (ICH2) USB controller USB-B> on uhci1
usb1: USB revision 1.0
uhub1: Intel UHCI root hub, class 9/0, rev 1.00/1.00, addr 1
uhub1: 2 ports with 2 removable, self powered
pcm0: <Intel ICH2 (82801BA)> port 0xd800-0xd8ff,0xdc40-0xdc7f irq 10 at device 31.5 on pci0
pcm0: [GIANT-LOCKED]
pcm0: <Analog Devices AD1885 AC97 Codec>
fdc0: <floppy drive controller> port 0x3f0-0x3f5,0x3f7 irq 6 drq 2 on acpi0
fdc0: [FAST]
fd0: <1440-KB 3.5" drive> on fdc0 drive 0
atkbd0: <Keyboard controller (i8042)> port 0x60,0x64 irq 1 on acpi0
atkbd0: <AT Keyboard> irq 1 on atkbd0
kbd0 at atkbd0
atkbd0: [GIANT-LOCKED]
psm0: <PS/2 Mouse> irq 12 on atkbd0
psm0: [GIANT-LOCKED]
psm0: model IntelliMouse, device ID 3
sio0: <16550A-compatible COM port> port 0x3f8-0x3ff irq 4 flags 0x10 on acpi0
sio0: type 16550A
sio1: <16550A-compatible COM port> port 0x2f8-0x2ff irq 3 on acpi0
sio1: type 16550A
pmtimer0 on isa0
orm0: <ISA Option ROMs> at iomem 0xc0000-0xc87ff,0xc8800-0xcbfff 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
Timecounter "TSC" frequency 996769408 Hz quality 800
```

Re: Odd performance problems after upgrade from 4.11 to 6.0–Stable

```
Timecounters tick every 1.000 msec
ad0: 38166MB <IC35L040AVER07 0 ER40A46A> at ata0-master UDMA100
ad1: 19092MB <IC35L020AVER07 0 ER20A46A> at ata0-slave UDMA100
acd0: CDRW <SONY CD-RW CRX160E/1.0g> at atal-master UDMA33
acd1: DMA limited to UDMA33, controller found non-ATA66 cable
acd1: DVDR <DVD-RW IDE1004/VER 0151> at atal-slave UDMA33
pcm0: measured ac97 link rate at 55947 Hz
Trying to mount root from ufs:/dev/ad0s1a
cd0 at atal bus 0 target 1 lun 0
cd0: <DVDRW IDE1004 0151> Removable CD-ROM SCSI-0 device
cd0: 33.000MB/s transfers
cd0: Attempt to query device size failed: NOT READY, Medium not present - tray closed
cd1 at atal bus 0 target 0 lun 0
cd1: <SONY CD-RW CRX160E 1.0g> Removable CD-ROM SCSI-0 device
cd1: 33.000MB/s transfers
cd1: Attempt to query device size failed: NOT READY, Medium not present
ipfw2 (+ip6) initialized, divert loadable, rule-based forwarding disabled, default to deny, logg
```

freebsd-stable@xxxxxxxxxxxxx mailing list

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

To unsubscribe, send any mail to "freebsd-stable-unsubscribe@xxxxxxxxxxxxx"