

# FreeBSD 6.1, crashes and a lack of vmcores

---

*Source:* <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/hackers/2006-05/msg00304.html>

---

- *From:* Matt Ruzicka <[matt@xxxxxxxx](mailto:matt@xxxxxxxx)>
  - *Date:* Fri, 26 May 2006 11:57:09 -0600 (MDT)
- 

For some time now we have been having a lot of trouble with one particular server which is part of a farm of six other largely identical servers. These servers run under extremely high load through a majority of the day and run a mix of postfix, MySQL (running as replication slaves) and custom filter software using MFS partitions. All seven servers are running on identical SuperMicro 6013E-i SuperServers with dual hyper-threading Xeon 2.80GHz CPU's with 2G of RAM. It is not all together uncommon for these machines to crash under extremely high load, but this one server in particular crashes much more frequently.

We started with memtest and CPU tests with no errors. As part of our troubleshooting we have replaced (or swapped out with the other servers) every piece of hardware in this box, replaced every cable and cord and moved to different switch and power ports. We've even changed physical locations in our data center. We have so far been unable resolve the more frequent crashes or move the increased instability to another server in an effort to find the cause. We've also disable hyper-threading in the bios and in FreeBSD on this machine since it sounds as if we might see other benefits from this. Also, as a stretch I've moved this box to using the ULE scheduler instead of the standard 4BSD. Really I'm starting to suspect it is haunted (or that I'm sleepdriving into work at night to foil my own progress).

These boxes traditionally run FreeBSD 4.11, but in a move of desperation we decided to take this particular machine up to FreeBSD 6.1 in an effort to rule out problems related to OS improvements and to ensure we are running the latest stable version of the different software pieces (and because it seems like the right move in the long term). (We install service software manually by the way, not from ports. MySQL we've installed from their binary distribution for 6.x.)

With the upgrade we are still receiving crashes at the same frequency and although the errors appear to report a bit differently they appear to be the same errors. Mostly a combination of "Fatal Trap 12" and "vm\_page\_fault" errors, though we have seen a couple "Sleeping thread owns a non-sleepable lock" errors.

The biggest frustration in this is that of the few dozen crashes we've had I've only been able to get one successful dump. All the other times I get the savecore error message:

```
kernel: kernel dumps on /dev/ad0s1b
kernel: Checking for core dump on /dev/ad0s1b...
kernel: unable to open bounds file, using 0
kernel: checking for kernel dump on device /dev/ad0s1b
kernel: mediasize = 4294967296
kernel: sectorsize = 512
kernel: magic mismatch on last dump header on /dev/ad0s1b
kernel: savecore: no dumps found
savecore: no dumps found
```

## FreeBSD 6.1, crashes and a lack of vmcores

Is there something I am missing to more reliably receive successful dumps? I have plenty of space on /var (22G) and my swap partition is 4G (with 2G of RAM).

The one successful dump returned the below gdb information. I've also included the non-commented bits of our kernel config at the very bottom.

If anyone has any suggestions on what this dump information indicates I would be very appreciative. Please let me know what other information I can furnish. If I can determine how to get another vmcore I'd be happy to send along another debug as well.

Thank you very much in advance.

Matt Ruzicka – Senior Systems Administrator  
Front Range Internet, Inc.  
matt@xxxxxxxx – (970) 212-0728

-----

```
[GDB will not be able to debug user-mode threads: /usr/lib/libthread_db.so: Undefined symbol
"ps_pglobal_lookup"]
GNU gdb 6.1.1 [FreeBSD]
Copyright 2004 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-marcel-freebsd".
```

```
Unread portion of the kernel message buffer:
vm_page_free: pindex(3255307648), busy(194), PG_BUSY(1), hold(-10260)
panic: vm_page_free: freeing busy page
cpuid = 0
Uptime: 18h43m26s
Dumping 2047 MB (2 chunks)
chunk 0: 1MB (159 pages) ... ok
chunk 1: 2047MB (524016 pages) 2031 2015 1999 1983 1967 1951 1935 1919 1903 1887 1871 1855 1839
1823 1807 1791 1775 1759 1743 1727 1711 1695 1679 1663 1647 1631 1615 1599 1583 1567 1551 1535
1519 1503 1487 1471 1455 1439 1423 1407 1391 1375 1359 1343 1327 1311 1295 1279 1263 1247 1231
1215 1199 1183 1167 1151 1135 1119 1103 1087 1071 1055 1039 1023 1007 991 975 959 943 927 911 895
879 863 847 831 815 799 783 767 751 735 719 703 687 671 655 639 623 607 591 575 559 543 527 511 495
479 463 447 431 415 399 383 367 351 335 319 303 287 271 255 239 223 207 191 175 159 143 127 111 95 79
63 47 31 15
```

```
#0 doadump () at pcpu.h:165
165 pcpu.h: No such file or directory.
in pcpu.h
(kgdb) where
#0 doadump () at pcpu.h:165
#1 0xc04b029d in boot (howto=260) at /u/frii/src/FreeBSD-6.1-RELEASE/sys/kern/kern_shutdown.c:402
#2 0xc04b05c5 in panic (fmt=0xc0600359 "vm_page_free: freeing busy page")
at /u/frii/src/FreeBSD-6.1-RELEASE/sys/kern/kern_shutdown.c:558
```

## FreeBSD 6.1, crashes and a lack of vmcores

```
#3 0xc05a2f45 in vm_page_free_toq (m=0xc207d7b0)
at /u/frii/src/FreeBSD-6.1-RELEASE/sys/vm/vm_page.c:1025
#4 0xc05a256d in vm_page_free (m=0xc207d7b0) at
/u/frii/src/FreeBSD-6.1-RELEASE/sys/vm/vm_page.c:403
#5 0xc059ff39 in vm_object_terminate (object=0xc878b4a4)
at /u/frii/src/FreeBSD-6.1-RELEASE/sys/vm/vm_object.c:631
#6 0xc059fe13 in vm_object_deallocate (object=0xc878b4a4)
at /u/frii/src/FreeBSD-6.1-RELEASE/sys/vm/vm_object.c:564
#7 0xc059c8fa in vm_map_entry_delete (map=0xc9f7e12c, entry=0xca3e2c38)
at /u/frii/src/FreeBSD-6.1-RELEASE/sys/vm/vm_map.c:2207
#8 0xc059cac7 in vm_map_delete (map=0xc9f7e12c, start=3335031932, end=3217031168)
at /u/frii/src/FreeBSD-6.1-RELEASE/sys/vm/vm_map.c:2300
#9 0xc059cb28 in vm_map_remove (map=0xc9f7e12c, start=0, end=3217031168)
at /u/frii/src/FreeBSD-6.1-RELEASE/sys/vm/vm_map.c:2319
#10 0xc0496fcd in exit1 (td=0xc9d93190, rv=0) at vm_map.h:211
#11 0xc04969b8 in sys_exit (td=0xc9d93190, uap=0x0)
at /u/frii/src/FreeBSD-6.1-RELEASE/sys/kern/kern_exit.c:97
#12 0xc05d8917 in syscall (frame=
{tf_fs = 59, tf_es = 59, tf_ds = -1079115717, tf_edi = -1077942712, tf_esi = -1077942820, tf_ebp =
-1077942876, tf_esp = -387965596, tf_ebx = 672734248, tf_edx = 10, tf_ecx = 672733680, tf_eax = 1,
tf_trapno = 12, tf_err = 2, tf_eip = 672673571, tf_cs = 51, tf_eflags = 646, tf_esp = -1077942904, tf_ss =
59}) at /u/frii/src/FreeBSD-6.1-RELEASE/sys/i386/i386/trap.c:981
#13 0xc05c58bf in Xint0x80_syscall () at /u/frii/src/FreeBSD-6.1-RELEASE/sys/i386/i386/exception.s:200
#14 0x00000033 in ?? ()
Previous frame inner to this frame (corrupt stack?)
(kgdb) up 2
#2 0xc04b05c5 in panic (fmt=0xc0600359 "vm_page_free: freeing busy page")
at /u/frii/src/FreeBSD-6.1-RELEASE/sys/kern/kern_shutdown.c:558
558 boot(bootopt);
(kgdb) p bootopt
$1 = 260
(kgdb) p *bootopt
Cannot access memory at address 0x104
(kgdb)
```

```
-----

machine i386
cpu I686_CPU
ident MAFILTER-NEW
makeoptions DEBUG=-g # Build kernel with gdb(1) debug symbols
options SCHED_ULE # ULE scheduler
options PREEMPTION # Enable kernel thread preemption
options INET # InterNETworking
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 NFSCLIENT # Network Filesystem Client
options PROCFS # Process filesystem (requires PSEUDOFS)
options PSEUDOFS # Pseudo-filesystem framework
```

## FreeBSD 6.1, crashes and a lack of vmcores

```
options COMPAT_43 # Compatible with BSD 4.3 [KEEP THIS!]
options COMPAT_FREEBSD4 # Compatible with FreeBSD4
options COMPAT_FREEBSD5 # Compatible with FreeBSD5
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
# output. Adds ~128k to driver.
options AHD_REG_PRETTY_PRINT # Print register bitfields in debug
# output. Adds ~215k to driver.
options ADAPTIVE_GIANT # Giant mutex is adaptive.
options SMP # Symmetric MultiProcessor Kernel
device apic # I/O APIC
device eisa
device pci
device ata
device atadisk # ATA disk drives
device atkbdc # AT keyboard controller
device atkbd # AT keyboard
device psm # PS/2 mouse
device kbdmux # keyboard multiplexer
device vga # VGA video card driver
device sc
device em # Intel PRO/1000 adapter Gigabit Ethernet Card
device miibus # MII bus support
device fxp # Intel EtherExpress PRO/100B (82557, 82558)
device loop # Network loopback
device random # Entropy device
device ether # Ethernet support
device tun # Packet tunnel.
device pty # Pseudo-ttys (telnet etc)
device md # Memory "disks"
```

---

freebsd-hackers@xxxxxxxxxxx mailing list

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

To unsubscribe, send any mail to "freebsd-hackers-unsubscribe@xxxxxxxxxxx"