

2.6.10 Kernel BUG at hugetlbpage:212 (x86_64 and i386)

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2005-02/0815.html>

From: Mr. Berkley Shands (*bshands_at_exeigy.com*)

Date: 02/03/05

Date: Thu, 03 Feb 2005 08:21:06 -0600

To: linux-kernel@vger.kernel.org

Reproducible BUG on 3GB hugetlbfs filesystem for opterons and xeon with either FC3 or RedHat ES3.0 and GCC 3.4.2. Details and code snippets in attachment. Executables to reproduce BUG are available on request.

berkley

On an 8GB dual cpu opteron (Tyan S2884) 2.6.10 kernel I can reproduce a crash within several minutes by creating/mapping/deleting/unmapping 512MB files using the code below. On a IA32 box (Xeon 2.4GHz, 3GB SuperMicro X5DA8) the crash is fatal (no /var/log/messages) and immediate. Executables for x86_64 and i386 from either FC3 or Redhat ES3.0 available on request. GCC is 3.4.2 on both machines and O/S releases.

```
Feb 2 11:28:49 noreaster kernel: ----- [cut here ] ----- [please bite here ] -----
Feb 2 11:28:49 noreaster kernel: Kernel BUG at hugetlbpage:212
Feb 2 11:28:49 noreaster kernel: invalid operand: 0000 [1] SMP
Feb 2 11:28:49 noreaster kernel: CPU 1
Feb 2 11:28:49 noreaster kernel: Modules linked in:
Feb 2 11:28:49 noreaster kernel: Pid: 15687, comm: DssiEPSearch Not tainted 2.6.10
Feb 2 11:28:49 noreaster kernel: RIP: 0010:[<ffffffff8011e3cb>]
<ffffffff8011e3cb>{unmap_hugepage_range+75}
Feb 2 11:28:49 noreaster kernel: RSP: 0018:00000100d7701dd8 EFLAGS: 00010206
Feb 2 11:28:49 noreaster kernel: RAX: 000001012eee98e0 RBX: 000001001000a0c0 RCX:
0000000061001000
Feb 2 11:28:49 noreaster kernel: RDX: 0000000061001000 RSI: 0000000041000000 RDI:
000001012eee98e0
Feb 2 11:28:49 noreaster kernel: RBP: 0000000041000000 R08: 0000000061200000 R09:
00000100d7701f10
Feb 2 11:28:49 noreaster kernel: R10: 0000000000000000 R11: 0000000000000202 R12:
0000000061001000
Feb 2 11:28:49 noreaster kernel: R13: 00000101f6ee0040 R14: 0000000041000000 R15: 000001012eee98e0
Feb 2 11:28:49 noreaster kernel: FS: 0000002a96b88080(0000) GS:ffff80611c00(0000)
```

Linux–Kernel: 2.6.10 Kernel BUG at hugetlbpage:212 (x86_64 and i386)

```
knlGS:0000000000000000
Feb 2 11:28:49 noreaster kernel: CS: 0010 DS: 0000 ES: 0000 CR0: 000000008005003b
Feb 2 11:28:49 noreaster kernel: CR2: 0000002a9632e970 CR3: 000000001184a000 CR4:
000000000000006e0
Feb 2 11:28:49 noreaster kernel: Process DssiEPSearch (pid: 15687, threadinfo 00000100d7700000, task
00000101215c63f0)
Feb 2 11:28:49 noreaster kernel: Stack: 0000000000000206 000001001000a0c0 0000000041000000
000001012eee9838
Feb 2 11:28:49 noreaster kernel: 00000101f6ee0040 0000000061200000 000001012eee98e0 ffffffff8015d800
Feb 2 11:28:49 noreaster kernel: 0000000000000000 ffffffff802bf59b
Feb 2 11:28:49 noreaster kernel: Call Trace:<fffffff8015d800>{unmap_vmas+320}
<fffffff802bf59b>{fpga_ioctl+2171}
Feb 2 11:28:49 noreaster kernel: <fffffff802c141b>{fpga_read+555} <fffffff80161a23>{do_munmap+547}
Feb 2 11:28:49 noreaster kernel: <fffffff801627a8>{sys_munmap+72}
<fffffff8010d216>{system_call+126}
Feb 2 11:28:49 noreaster kernel:
Feb 2 11:28:49 noreaster kernel:
Feb 2 11:28:49 noreaster kernel: Code: 0f 0b 33 1f 49 80 ff ff ff ff d4 00 4c 89 f5 4d 39 e6 66 66
Feb 2 11:28:49 noreaster kernel: RIP <fffffff8011e3cb>{unmap_hugepage_range+75} RSP
<00000100d7701dd8>
Feb 2 11:32:32 noreaster kernel: Linux version 2.6.10 (root@noreaster) (gcc version 3.4.2) #47 SMP Wed
Feb 2 08:01:40 CST 2005
```

```
// code to access a huge tlb filesystem
```

```
#if defined(__linux__) && defined(__x86_64__)
HugePageFileName_ = new char[64];
::sprintf(HugePageFileName_, "/mnt/huge/Silo_XXXXXX");
::mkstemp(HugePageFileName_); // randomize this name
Mmap_Fd_ = ::open(HugePageFileName_, O_CREAT | O_RDWR | O_LARGEFILE | O_TRUNC, 0755);
if (Mmap_Fd_ != -1)
{
    LONG64 MyBig = Big * sizeof(LONG64) + ((2UL * 1024UL * 1024UL) - 1);
    MyBig &= ~((2UL * 1024UL * 1024UL) - 1);
    BigBlock0_ = (LONG64 *) ::mmap(NULL, MyBig,
        (PROT_READ | PROT_WRITE),
        MAP_SHARED, Mmap_Fd_, 0);
    if (BigBlock0_ == MAP_FAILED || !BigBlock0_)
    {
        ::perror("mmap failed for huge pages");
        ::cerr << "Asked for ";
        printHRNumber(MyBig, cerr);
        ::cerr << endl;
        ::close(Mmap_Fd_);
        ::unlink(HugePageFileName_);
        delete [] HugePageFileName_;
        HugePageFileName_ = NULL;
        Mmap_Fd_ = -1;
        BigBlock0_ = NULL;
        BigBlock_ = NULL;
    }
}
```

Linux–Kernel: 2.6.10 Kernel BUG at hugetlbpage:212 (x86_64 and i386)

```
else
{
BigBlock_ = BigBlock0_;
BigSize_ = MyBig; // remember this size
::unlink(HugePageFileName_); // pre–delete this file so segfaults
// do not leave huge pages tied up.
}
}
else
{
if (Debug_)
{
::perror("Sorry, no HUGE pages today");
}
}
}
#endif
```

```
// snip from /etc/rc.local
```

```
#
# setup the huge file system space pages are 2MB each
#
if [-f /proc/fpga0]; then
echo "Allocating huge file system"
echo 1536 > /proc/sys/vm/nr_hugepages
mount none /mnt/huge -t hugetlbfs -o size=3G,mode=0777
fi
```

–

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/>