

## Re: Lockless uidinfo.

---

*Source:* <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/arch/2007-08/msg00043.html>

---

- *From:* Pawel Jakub Dawidek <[pjd@xxxxxxxxxxxx](mailto:pjd@xxxxxxxxxxxx)>
  - *Date:* Sun, 19 Aug 2007 09:57:50 +0200
- 

On Sat, Aug 18, 2007 at 04:35:42PM -0700, Jeff Roberson wrote:

On Sun, 19 Aug 2007, Pawel Jakub Dawidek wrote:

Ok, after implementing `atomic_fetchadd_long()` on amd64, we get additional 6% of performance improvement:

```
x ./uidinfo_lockfree.txt (atomic_cmpset_long loop)
+ ./uidinfo_waitfree.txt (atomic_fetchadd_long)
```

```
+-----+
|
+|
|
+|
|x xx xx
+ ++|
| |__MA__|
|AM|
+-----+
```

```
N Min Max Median Avg Stddev
x 5 1561566 1575987 1568964 1569767 5853.1399
+ 5 1662362 1665936 1665810 1664881.8 1541.2693
Difference at 95.0% confidence
95114.8 +/- 6241.96
6.05917% +/- 0.397636%
(Student's t, pooled s = 4279.88)
```

How does this effect the single-threaded performance? Do you attribute this to atomic fetchadd being cheaper than atomic cmpset? What is your processor?

CPU: Intel(R) Xeon(R) CPU E5310 @ 1.60GHz (1597.65-MHz  
K8-class CPU)

Origin = "GenuineIntel" Id = 0x6f7 Stepping = 7

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

Features2=0x4e33d<SSE3,RSVD2,MON,DS\_CPL,VMX,TM2,SSSE3,CX16,xTPR,PDCM,DCA>

AMD Features=0x20100800<SYSCALL,NX,LM>

Re: Lockless uidinfo.

AMD Features2=0x1<LAHF>  
Cores per package: 4

Ok, I changed the code to something like this:

```
long old;
int diff, loops;

atomic_add_int(&uidinfo_cnt1, 1);
if (diff > 0) {
    loops = 0;
    do {
        loops++;
        old = uip->ui_sbsize;
        if (old + diff > max)
            return (0);
    } while (atomic_cmpset_long(&uip->ui_sbsize, old, old + diff) == 0);
    if (loops > 1)
        atomic_add_int(&uidinfo_cnt2, loops);
    } else {
        atomic_add_long(&uip->ui_sbsize, (long)diff);
    }
}
```

This allows me to see how many additional loops I do, because with lock-free version we still can have contention and loop, that's why wait-free version is superior.

Actually I was a bit surprised with the results:

```
debug.uidinfo.cnt1: 88746008
debug.uidinfo.cnt2: 31296304
```

(Running 8 processes.)

Which means, because of contention, we do 31296304 additional atomic operations, which is about 30% more.

—

Pawel Jakub Dawidek <http://www.wheel.pl>  
pjd@xxxxxxxxxxx <http://www.FreeBSD.org>  
FreeBSD committer Am I Evil? Yes, I Am!

**Attachment:** [pgpJotQKWNQGD.pgp](#)

*Description:* PGP signature

Re: Lockless uidinfo.