

Re: ten thousand small processes

Source: <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/performance/2003-06/0098.html>

From: D. J. Bernstein (*djb_at_cr.yp.to*)

Date: 06/26/03

Date: 26 Jun 2003 02:50:29 -0000

To: freebsd-performance@freebsd.org

Jon Mini writes:

> *I'm sorry, but you are way off here. First of all, caches are *much
> larger* than the size of the processes you are talking about.*

I'm sorry, but you are being misled by a naive model of CPU performance. On a typical Pentium in our department, the following program becomes three times faster when SPACING is changed from 4096 to 128:

```
#define SPACING 4096
char data[8 * SPACING];
main()
{
    int i;
    for (i = 0; i < 10000000; ++i) {
        data[0] = data[SPACING];
        data[2 * SPACING] = data[3 * SPACING];
        data[4 * SPACING] = data[5 * SPACING];
        data[6 * SPACING] = data[7 * SPACING];
    }
}
```

> *From an asm programmer's perspective, when FreeBSD decides to spread a small program's variables between*

- * the beginning of a data page,
- * the beginning of a bss page,
- * the beginning of a malloc mmap page,
- * the beginning of a heap page,
- * the beginning of the next heap page,
- * the beginning of yet another heap page,

et cetera, it is actively trying (with varying degrees of success) to damage cache performance in exactly the same way that this program does.

—D. J. Bernstein, Associate Professor, Department of Mathematics, Statistics, and Computer Science, University of Illinois at Chicago

freebsd-performance: Re: ten thousand small processes

freebsd-performance@freebsd.org mailing list

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

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