

freebsd-current: Re: low(er) disk performance with sched\_4bsd then with sched\_ule

## Re: low(er) disk performance with sched\_4bsd then with sched\_ule

**Source:** <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/current/2005-09/0610.html>

---

**From:** Andrew Gallatin ([gallatin\\_at\\_cs.duke.edu](mailto:gallatin_at_cs.duke.edu))

**Date:** 09/17/05

Date: Sat, 17 Sep 2005 17:33:38 -0400 (EDT)

To: [joseph.koshy@gmail.com](mailto:joseph.koshy@gmail.com)

Oliver Lehmann writes:

> *Joseph Koshy wrote:*

>

> > *ol> Wow, that update to BETA4 did the trick! While running*

> > *ol> SCHED\_4BSD:*

> >

> > *Fantastic! What is the profile like with the new 4BSD kernel?*

>

> <http://pofo.de/tmp/gprof.4bsd.3>

I don't know the disk codepath very well, but the samples look a little suspect. We're copying a lot of data into and out of the kernel, so I would expect the majority of non disk wait time would be spent simply copying out the zero-filled pages, and copying them back in (AFAIK, dd uses read/write). Where is the time spent in read, write, uiomove, bcopy?

What about ionode allocations, etc? And why do things like `g_bsd_modify` and `g_bsd_ioctl` rank so high? Aren't those only used when dealing with disklabels?

BTW, I *\*love\** that we've got access to the hw counters, and an easy way to do low-overhead profiling of the kernel.

Drew

---

freebsd-current@freebsd.org mailing list

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

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