

## Re: headsup: swap\_pager.c

**Source:** <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/arch/2003-08/0003.html>

---

**From:** David Schultz (*das\_at\_FreeBSD.ORG*)

**Date:** 08/01/03

Date: Fri, 1 Aug 2003 01:23:36 -0700  
To: Poul-Henning Kamp <phk@phk.freebsd.dk>

On Thu, Jul 31, 2003, Poul-Henning Kamp wrote:

- > *People who do not adjust NSWAPDEV to match their configuration spend*
- > *a lot of RAM on holding the radix bitmap which is 3/4 empty:*
- >
- > *critter phk> sysctl kern.malloc | grep -i swap*
- > *SWAP 2 753K 753K 2 64*
- > *critter phk> pstat -s*
- > *Device 1K-blocks Used Avail Capacity*
- > */dev/ad0s1b 1321312 0 1321312 0%*
- >
- > *This is because the radix bitmap gets allocated for the full*
- > *stripe width, but only one quarter of it is actually used.*

Static striping has many annoying limitations, but the memory footprint isn't so significant. If people cared enough, we could save a factor of 2 by fixing the bugs that were introduced when `daddr_t` became 64 bits, and a factor of 8 by tracking swap space in `PAGE_SIZED` quantities instead of `DEV_BSIZED` quantities. By comparison, more clever striping strategies probably have a higher memory footprint, but have better flexibility and performance.

But since you just ripped out the debugging stuff for the radix tree bitmap, I'm guessing that your plan is to replace the bitmap with something else within the next few days anyway...

- > *Another thing is that striping does not belong in the swap\_pager in*
- > *the first place, we have CCD and similar pieces of code for that.*
- >
- >
- > *The new layout will allocate the devices sequentially in the "virtual*
- > *swapdevice" and allocation will (for now) be round robin between*
- > *the devices. Later I would like to implement load based device*
- > *selection (ie: pick the device with the lowest recent average access*
- > *time).*
- >
- > *If disk striping is necessary for the disks used for swap-space,*
- > *do it with ccd(4) and call swapon on that.*

freebsd-arch: Re: headsup: swap\_pager.c

Can you please describe how the interface between ccd(4) and the swap subsystem will work? I like this idea, but it isn't immediately obvious how you plan to implement striping based on load and dynamic addition and removal of swap space.

---

freebsd-arch@freebsd.org mailing list

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

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