

Re: 20TB Storage System (fsck????)

Source: <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/performance/2003-09/0024.html>

From: Chuck Swiger (cswiger_at_mac.com)

Date: 09/03/03

Date: Wed, 03 Sep 2003 14:37:50 -0400

Max Clark wrote:

- > *Ohh, that's an interesting snag. I was under the impression that 5.x w/ PAE*
- > *could address more than 4GB of Ram.*

It can. PAE lets the hardware address more than 4GB of RAM, but that doesn't change how much memory you can give to any one process: a 32-bit process still has a 32-bit virtual address space.

- > *If fsck requires 700K for each 1GB of Disk, we are talking about 7GB of Ram*
- > *for 10TB of disk. Is this correct? Will PAE not function correctly to give*
- > *me 8GB of Ram? To check 10TB of disk?*

Another thread suggests that the maximum amount of memory actually available for a 32-bit process to use under FreeBSD is a little less than 3 GB.

- > *Is there anyway to bypass this requirement and split fsck into smaller*
- > *chunks? Being able to fsck my disk is kinda important.*

Sure. Create multiple filesystems rather than just one, and use symlinks to make the directory namespace fit your needs.

I don't know enough about your tasks to give you really specific advice, but I'm wary of the write-performance hit from putting too many drives wide in a RAID-5 (or -5,0) configuration. If you can split up your data by role or typical access pattern, you might well be able to identify some chunks that will be read-mostly (and RAID-5,0 is a good fit) and others that will be read-write or even write-mostly (and thus should be on -1,0). You can also tune other things like blocksize, # of inodes, and so forth more appropriately for each filesystem.

--

-Chuck

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"