

strange performance dip shown by iozone

Source: <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/performance/2004-02/0043.html>

mi+mx_at_aldan.algebra.com

Date: 02/18/04

To: fs@FreeBSD.org, performance@FreeBSD.org

Date: Wed, 18 Feb 2004 17:29:06 -0500

Hello!

I'm trying to tune the amrd-based RAID5 and have made several iozone runs on the array and -- for comparison -- on the single disk connected to the Serial ATA controller directly.

The RAID-based FS was newfs-ed with ``-b 65536'', as it is intended to store very large files. The single-disk FS was newfs-ed with defaults.

No softupdates were enabled on either, since those seem to degrade iozone results slightly (iozone reads/writes a single file anyway).

The filesystems displayed different performance (reads are better with RAID, writes -- with the single disk), but both have shown a notable dip in writing (and re-writing) speed when iozone used the record lengths of 128 and 256. Can someone explain that? Is that a known fact? How can that be avoided?

The machine is an amd64 running a fresh -current. The disks are 200Gb SATAs. RAID5 consists of 6 of them.

Thanks!

-mi

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"