

Re: ZFS performance weirdness on a a Netra X1

Source: <http://unix.derkeiler.com/Newsgroups/comp.unix.solaris/2006-08/msg01817.html>

- *From:* "Daniel Rock" <v200633@xxxxxxxxxxxx>
 - *Date:* Fri, 18 Aug 2006 16:43:27 +0000 (UTC)
-

Daniel Rock <v200633@xxxxxxxxxxxx> wrote:

 Hmm, tomorrow I might be able to do some tests on old hardware as well (E3500 with 400 MHz CPUs).

Ok, I have done some tests on this ancient brick with mixed results. It is a E3500 with 4 x 400 MHz. I have turned off 3 of its CPUs (via psradm) to get similar results. Disks were FC attached A5200 disks (since ZFS doesn't reset the write cache if you destroy a pool the UFS tests did also benefit from the still enabled write cache).

I achieved with ZFS a write throughput
dd if=/dev/zero of=/pool/bigfile bs=1024k count=1024
sync
an average throughput of ~15 MB/s
But during the test the system was almost unusable (sluggish) with a kernel time of up to 99%.

With UFS in contrast I achieved the same throughput but the system kept being responsive during the test (kernel time at ~35%)

So the CPU is likely too slow for ZFS. Turning checksums off didn't make any significant difference.

--
Daniel

.