

freebsd-questions: Re: What do you use?

Re: What do you use?

Source: <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/questions/2004-01/0115.html>

From: Scott W (wegster_at_mindcore.net)

Date: 01/02/04

Date: Fri, 02 Jan 2004 00:42:38 -0500

To: Scott Mitchell <scott+freebsd@fishballoon.org>

Scott Mitchell wrote:

>On Thu, Jan 01, 2004 at 05:59:18PM -0500, Scott W wrote:

>

>

>>Scott Mitchell wrote:

>>

>>

>>>There no particular reason for an ATA RAID to be slower than SCSI, assuming

>>>similar disks in each. 10krpm 'server class' ATA disks are available these

>>>days, although I don't know that anyone has done a 15krpm one yet.

>>>

>>>

>>>

>>Does SATA have tagged queing? (I don't know offhand if it does...?)

>>

>>

>

>No idea. I haven't used any SATA stuff yet. Would be surprised if it

>didn't though.

>

It doesn't look like it does, although it's supposedly on the way with SATA-II, but see below...

>

>

>>I can guarantee modern SCSI throughput is superior to any of the SATA

>>drives I've seen to date. Several of the 'hardware sites' (I think

>>Tomshardware did a writeup on this or anadtech among others) agree with

>>this statement as well. ATA specs tend to exaggerate their capabilities

>>even worse than SCSI specs do- burst speeds are all fine and dandy, but

>>not realistic at all in the real world. Meaning basically in short I

>>wouldn't choose SATA over SCSI for a production server of any kind where

>>speed was an issue. ATA has gotten better by far than it was

>>speed-wise, and I'd be OK with it on a personal workstation for any

>>purpose, but it's still playing catchup.

>>

Re: What do you use?

freebsd-questions: Re: What do you use?

>>

>

>Well, I won't disagree with any of that – ATA drives have been aimed at the
>consumer market, which wants big & cheap in preference to performance &
>reliability. On the other hand, you've got drives like the WD 'Raptor'
>that seem to be mechanically what you expect in a server-class SCSI drive,
>but with a SATA interface, so the gap isn't that big anymore. I stand by
>my statement above though – there's no reason an ATA RAID implementation
>_has_ to be slower than SCSI, provided that the disks are up to it. In
>practice though, suitable disks are only just now starting to appear.

>

I did some digging on this, mainly because I remember reading a specific article. The WD Raptor really IS pretty impressive, I've got to admit. Some benchmarks for single disks compared to an Atlas IV 10K RPM SCSI disk at Tom's Hardware:

<http://www20.tomshardware.com/storage/20030501/wd360-01.html>

So, it looks like things may actually become 'more interesting' over the next year or two basically. I used to like WD drives quite a bit, but have also had a fair number of them die within a few years (as well as other IDE/ATA desktop drives), be interesting to see if they can offer the warranties and constant uptime of SCSI drives over time...(some of the SATA drives only have 1yr warranties, where 5yrs is pretty standard for SCSI..)

Pricing for the Raptor 36GB versus U320 SCSI 36GB drives is pretty comparable (raptors ~\$130, Seagate SCA 10K U320 SCSI for ~\$135 via pricewatch), but the larger SATA drives are still only 7200RPM, where you can get 15K RPM SCSI drives.

From what I can tell, one of the current problems aside from spindle speeds for SATA is that they are exactly that– serial connections that can't be chained, so effectively need one channel per disk. For home workstations, that's fine with two channels, and better yet with a cheap controller that will do hardware disk striping, but doesn't do much compared to SCSI and fiber channel RAID enclosures holding 10–14 drives each. I've seen 4 channel SATA controllers, which can just barely do RAID–5 reasonably (2 data, 1 parity, 1 hot spare, effective, not physical), and it does look like 8 channel controllers are available, possibly a bit less than their SCSI RAID counterparts. I was pretty impressed– check out the CPU utilization during the benchmarks...ignore the other ATA and SATA drives, but look at SCSI and the WD Raptor– ~2% CPU utilization while pushing a _lot_ of data (opposed to 20%+ for other ATA drives..)

Anyways, it will really be interesting to see if the SATA drives can last long-term and then compare 'em against 15–20K RPM SCSI RAID or if Serial SCSI is in a year or so– one of the interesting points is that SCSI using a shared bus, has an effective max of 320MB/second for each bus, which in turn can be up to 15 drives on each, where SATA, being serial, has it's limit per channel, but containing a single drive. If

Re: What do you use?

freebsd-questions: Re: What do you use?

the SATA drives can be made in reasonable capacities at 15K (or the next generation speeds) with a controller capable of handling 15 disks (or more), that could potentially become a real threat to SCSI.

>
>
>> *Write performance is awful locally, or over NFS? NFS isn't exactly a speed demon.*
>>
>>
>
> *Locally – with RAID-5 you have the extra overhead of calculating the parity for the block you just wrote, then writing that out to an other drive.*
> *According to the Vinum docs, write performance is ~25% of read, although how much of that is Vinum and how much is just the nature of RAID-5 I'm not sure.*
>
Yep– part of the problem with software RAID. On a non-heavily loaded server it's not an issue, as you might as well use the CPU cycles, but on a loaded database server, not so great. FWIW, on a Sun box with a fiber channel array with an integrated Mylex based raid controller, reads and writes are within 1%, at least given the same quick and dirty test (using dd creating 32k blocks writing out a 2GB file). Part of that 25% is due to using the system CPU(s) versus hardware to calculate the parity data, plus handling the individual writes to each individual disk, but the majority of the difference is probably due to the onboard cache on the RAID controller. Older Sun boxes used to have an optional NVRAM module – without the module it wasn't worth using as a file server, it made that much difference...

Scott

>
>
>> *No comment on the unlimited budget as everyone at work just got (another) 'mandatory pay reduction'...but I do rememeber and miss those,*
>> *\$^#&*(*
>> *;-)*
>>
>>
>
> *:-)*
>
> *Scott*
>
>
>

freebsd-questions@freebsd.org mailing list
<http://lists.freebsd.org/mailman/listinfo/freebsd-questions>

Re: What do you use?

freebsd-questions: Re: What do you use?

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