Re: 306GB drives!

Source: http://unix.derkeiler.com/Newsgroups/comp.os.vms/2003−08/1117.html

---

From: Bob Lail (*Robert.Lail_at_hp.com*)
Date: 08/19/03

Date: Tue, 19 Aug 2003 18:36:49 GMT

"Bill Todd" <billtodd@metrocast.net> wrote in message
news:VaidndOsz4WOddyiXTWJkw@metrocast.net...
>
> *"Bob Lail" <Robert.Lail@hp.com> wrote in message*
> *news:btb0b.2619$3s2.1400@news.cpqcorp.net...*
>
> *...*
>
> *> > Read the paper.*
>
> *Ok.*
>
> *It will enlighten you to the real differences between*
> *> ATA and SCSI disk drives.*
>
> *Though it was a good overview, I'm afraid that there really wasn't*
anything
> *surprising in it.*
>
> *> > They are VERY different*
>
> *Of course they are: the question is the degree to which the differences*
> *matter. There are certainly cases where they are absolutely critical to*
> *OEMs (e.g., those who require modifications to the disk's firmware, or*
> *specialized sector sizes), but this is not true for typical customer use*
> *unless there's a need, say, to connect individual drives (rather than*
> *multi−drive cabinets) to multiple hosts.*
>
> *and no one should be*
> *> advocating replacing SCSI drives with ATA drives in Mission Critical*
> *> applications unless they know and understand the risks they will be*
taking
> *> with their data.*
>
> *Since I have a pretty good understanding of those risks, I'm fairly*
> *comfortable explaining them to others. They primarily boil down to*
> *potential performance degradation and premature failure if careful*

attention
> *is not paid to the operating environment.*
>
> *The article makes it clear that the increased cost of SCSI disks arises*
> *primarily from their significantly higher unit performance (more complex*
> *interface, higher rotational speeds and faster seeks, plus mechanical and*
> *electrical engineering that provides somewhat higher MTBF and lower retry*
> *frequencies despite these higher stresses) − in combination, of course,*
with
> *the consequences of spreading development costs over far smaller volumes.*
> *By contrast, ATA drives benefit from trickle−down engineering advances and*
> *less demanding operating characteristics.*
>
> *So while it's true that a 7200 rpm ATA drive has only about half the*
> *random−access performance of a 15Krpm SCSI/FC drive (which costs close to*
> *10x as much for a given storage capacity: it's only the older, slower 10K*
> *rpm SCSI generation that gets the difference down close to 3x), if there's*
> *enough parallelism in the workload you can simply spread the high−end*
> *drive's data over two ATA drives and obtain similar performance (assuming*
> *that the ATA drives are mounted so as to minimize mechanical*
seek−vibration
> *coupling, and modulo any additional differences due to the possible*
absence
> *of tagged−queuing facilities in the ATA drives, but the latter becomes*
> *significant only if the load is allowed to increase to the point where*
> *non−negligible queues actually start to form at the disks). And as shared*
> *virtualization technology increasingly spreads unrelated activities over*
> *many shared spindles, situations in which no significant parallelism is*
seen
> *at the disks become increasingly rare.*
>
> *There's no question that high−end drives are more reliable on an*
individual
> *basis, and that this advantage is even greater (by a factor of 3 − 4,*
> *according to the paper − which was a useful quantitative nugget of*
> *information that I hadn't previously been familiar with, including the*
> *compensating effect of reducing the platter count which once again can*
> *operate to reduce any ATA handicap) when adjusted to compensate for their*
> *higher nominal power−on and seek duty cycles, but their commodity*
> *competition is not one−for−one here any more than it is when performance*
is
> *important. A pair of mirrored commodity ATA drives provides comparable*
> *(read) performance to and far better (system) MTBF − even in 24/7*
> *operation − than the best single high−end drive you can buy, plus far*
> *greater storage capacity, far lower cost, and comparable overall*
> *power/cooling requirements (which allows each individual commodity drive*
to
> *operate at lower temperature − a reliability advantage over higher−powered*
> *SCSI/FC drives). The high−end drive's only real advantages are smaller*
rack
> *footprint (if you don't need the far greater storage capacity that the*

> *paired−ATA option provides), faster strictly−serial access, faster write*
> *access (where mirroring doesn't help the performance of the ATA pair),*
less
> *sensitivity to mechanical seek−coupling (though since ATA drives*
generate
> *less seek vibration in the first place, due to their less−aggressive head*
> *movements, that helps compensate), more graceful performance degradation*
> *under high loads than ATA drives that don't include internal queuing*
> *optimization, and less frequent replacement.*
>
> *While the paper is a good one, one should remember that its authors work*
for
> *a company with a major incentive to justify continued use of its*
high−margin
> *storage products. So it's hardly surprising that they choose to highlight*
> *the advantages of their high−end products rather than explain how*
commodity
> *products might effectively replace them in many uses.*
>
> *− bill*
>
>
>

Bill

I don't disagree with most of your comments, except maybe that 2 * ATA drive can match the sustained performance of a single 15K SCSI drive. I would be more inclined to go for a 3 to one ratio and some of that is dependent on the ATA controllers abilities to do the striping. Your 10x price analogy is coming down as it is becoming increasingly difficult to take more cost out the ATA drives while process improvements an increasing quantities in SCSI drives are starting to bring their costs down, at least on parallel SCSI drives. Fibre Channel drives are still to new, with not enough vendor competition and adoption to drive the cost down. Of course with only a few, three I believe (Seagate, Maxtor, and Fujitsu), SCSI drive vendors left in the marketplace there is a limit to that. ATA, I believe adds Hitachi and Toshiba to that list..

As you state virtualization technologies will make for some very interesting opportunities for SATA drives in the near future. As for size the next generation will be based upon 2/1/2" platters in both ATA and SCSI reducing the footprint and from what I have seen SATA and SAS drives will be exactly the same footprint in 2/1/2" technology.

SATA (V2) will address the lack of tag queuing with today's parallel ATA and first generation SATA drives so that will help the controller side of the ATA equation.

SATA (Serial ATA) and SAS (Serial Attached SCSI) share the same physical infrastructure allowing for both to exist at the same time on a physical

serial bus. There is a small difference in the physical connector that will prevent mix−ups. You will be able to mix and match SATA and SAS drives to their usage pattern in the same disk array. Use SATA when the usage pattern is low and does not require high performance and SAS when you need the performance and reliability of the more expensive SAS devices. The controller will adapt auto−magically to the required protocol.

As for you last comment, yep it would be nice to get an analysis from the ATA camp to compare.

As the Chinese curse goes, "May You Live in Interesting Times"

\Bob Lail

−−
Robert G. Lail
Senior Solution Architect
Corporate Accounts
Hewlett−Packard Company
Voice: 603.424.6272
Email: RobertDOTLail@hpDOTcom
(replace the DOT with ".")