STAR Watch

Statewide Technology Assistance Resources Project

A publication of the Western New York Law Center, Inc.

Volume 9 Issue 2 March/April 2005

Hard Drive Capacity: How much do you really need?

Computer geeks will always advise that there is no such thing as a computer hard drive that is too big or too fast. While those among us who use computers only for word processing or spreadsheets may be satisfied with a hard drive capacity of 20 – 40 Gigabytes, there are some of us who need much, much more. For instance, anyone who has ever edited a video will tell you that 3.16 Megabytes of disk storage is required to store 1 second of raw video. That translates into approximately 12 Gigabytes of hard drive storage for every hour of raw video, and doesn't include any work/temp files or the final product. Many of us with digital still cameras are also running out of space. While those new high-resolution cameras produce excellent quality images, the amount of storage space for each picture can exceed 2 megabytes per picture.

With the latest products announced by hard drive manufacturers, it is now possible to purchase computer hard drives that have a capacity of up to 400 Gigabytes. With that amount of storage, users can store over 200,000 digital photos or 33 hours of raw video. That should take a while to fill up.

Many people would like to upgrade the hard drive on their computer or purchase a new computer with a lot more storage capacity, but become intimidated by all of the acronyms and buzzwords. It isn't that difficult. There are only a few issues that need to be addressed.

SCSI vs SATA vs PATA

SCSI, SATA, and PATA describe how the hard drive is physically connected into the computer. Each of these interfaces have advantages and disadvantages that users should be aware of. None of these interface types are

In this issue…

compatible with any of the others. If you are upgrading a computer, you must install the correct type of hard drive, or be prepared to purchase and install a new interface card to connect the hard drive to the computer.

Typically, SCSI interface hard drives are most commonly found in file servers and a few very high performance workstations. Their performance is somewhat better than PATA or SATA drives, but not their price. A SCSI hard drive costs 4 to 8 times as much as a PATA or SATA drive of the same capacity. The typical computer user cannot justify the high cost of these drives for a desktop computer.

PATA, or Parallel ATA is also referred to as IDE or EIDE. It is currently the most prevalent hard drive type in computer workstations today. These types of drives are connected into the computer by a wide, flat, ribbon cable. Some critics argue that if these ribbon cables are improperly routed, they can cause cooling problems by blocking airflow to critical components inside the computer case.

SATA, or Serial ATA interface hard drives have been available for at least 2 years, and expected to outsell PATA drives by the middle of this year. SATA drives are connected to the computer by a thin interface cable that cannot block airflow inside the case.

A final word about PATA vs SATA drives: If you compare the technical specs of any manufacturer's PATA and SATA drives, it is very clear that the physical storage device at the end of the data cable is the same. Only the electronic interface is different. According to the published test results, SATA drives are slightly faster.

If the hard drive is going to go into an existing computer, it would be much better to use the same type of drive. While it is possible to mix and match drive types or use different drive types, it could get very complicated. When upgrading storage on an existing computer, stick to the type of drive that is already in the computer. This is not an issue with new computer purchases.

RAID (Redundant Array of Independent Disks)

While many people have gotten the impression that RAID is just another interface type like SCSI, PATA, and SATA, it isn't. RAID provides a means to connect multiple hard drives together to achieve greater storage capacity, better performance, greater storage reliability, or any combination of the three. As such, the subject of RAID is outside the scope of this article, except to say that there are RAID systems for SCSI, PATA and SATA hard drives.

A Word About "Drive Speeds"...

There is a lot of advertising out there that describes the speed of PATA and

SATA hard drives as "ATA-100", "ATA-133", or "ATA-150". While the numeric portion of the description does accurately describe the maximum speed (in megabytes/sec.) of the data transfer from the cache buffer to RAM (see Cache Buffer Size below), it only measures a very small part of the work that the hard drives does. The drive must retrieve data physically from the disk and place it into the cache buffer, or write the data found in the cache buffer onto the disk. The ATA speed measurement never considers the most time-consuming part of the job.

The true measurement of overall drive performance is "maximum sustainable data rate". This measurement describes the average number of megabytes of data that the drive can move from disk to RAM or move from RAM to disk over an extended period of time. Manufacturers will quote a number that does not take into consideration processor performance or CPU load caused by applications, which can greatly affect the actual performance of the drive in the real world. Real world results may be quite different.

At home, I have two computers. In each of those computers, are PATA hard drives from the same manufacturer with identical performance specs (7200 RPM, ATA-133). As ATA-133 drives, one would assume that the drives will be able to move data at 133 megabytes/sec. But, in the fine print, the manufacturer states that the maximum sustained data rate is 60 megabytes/sec. When I run speed tests on the drives, one computer had a sustained data rate of 29 – 31 megabytes/sec. The other achieved 49 – 50 megabytes/sec. The computer with the "slower" results is a 1.3GHz Pentium 4 with 256 MB RAM. The computer with the "faster" hard drive has a 3.4 GHz 64-bit processor with 1 GB of RAM.

Real world performance is greatly affected by the rest of the computer hardware. Don't obsess over the "ATA" speed. Even though the manufacturers make a big deal about it, it doesn't have that much affect on real world performance.

Rotational Speed

Faster is better. Faster will make a noticeable difference. As the RPMs on a hard drive increase, the time required to move the desired data under the read/write head decreases. Once located, faster RPMs mean faster physical reads and writes of

data. Most hard drives sold for use in desktop computers today are 7200 RPM. This is a significant improvement over the 5400 RPM drives that were the standard a few years ago. That increase in speed translates into a 25% reduction in the time required to physically access data. The 5400 RPM hard drives are still available for what might seem like a fantastic price, but the reality is that their lack of performance more than negates the price advantage.

While there are also hard drives that spin at 10,000 and 15,000 RPM, most are SCSI, only a few 10K drives are SATA, and none are PATA. There are very few SCSI drives spinning at 10K or 15K with at capacity greater than 147 Gigabytes. Currently, the largest capacity 10K SATA drive is 74 Gigabytes. These drives are mostly intended for use in file servers where the demands can justify the additional cost. For desktop computers, stick with 7200 RPM drives.

Cache Buffer Size

Once the data has been physically retrieved from the hard disk, it must be moved to a location where the application program can use it. Conversely, if an application is writing data, it must be moved to a location where it will be ready to be physically written to disk. That brings up an issue: As the computer gets busier, more and more data requests must be processed. It becomes more and more of a problem to stream the data to/from the computer application to the disk in a single seamless operation. The cache buffer on the hard drive allows the process of transferring data to/from the hard drive to occur in 2 distinct steps, which reduces the number of issues that must be controlled at any point in time.

Bigger is better and becomes more important as the speed of the processor increases. The greater the cache buffer size, the greater the possibility that the drive has a place to put data after retrieving it from disk or that it has data that can be written to disk. Most hard drives sold today have an 8 megabyte disk cache. While it is still possible to find some drives with a 2 megabyte cache for a lower price, the loss in performance would negate any cost savings.

One vendor, Maxtor, has increased the cache buffer size on some of its drives to 16 megabytes. Depending on the disk I/O load caused by the applications running on the computer, disk performance could be improved. It definitely wouldn't slow it down.

Native Command Queuing (NCQ)

One of the most innovative features in this new offering of hard drives is Native Command Queuing. Without NCQ, hard drives process requests in the exact order that they were requested. With NCQ, requests are handled in an order that is most efficient for the hard drive. This could increase hard drive performance in peak load situations, which could speed up applications running on the computer.

So how does it work? Consider the following situation: The hard drive has received commands to read data alternately from the inner then outer portions of the disk. Without NCQ, the drive would dutifully read, then move the heads to the far side of the disk, read a record, move all the way back, etc. The read/write heads are spending a great deal of time moving back and forth to get to the next record. With NCQ invoked, the drive would complete all of the reads in the general location where the read/write heads currently are stationed, then move to the far side of the disk and do the rest of the reads. In this example, the read/write heads would only travel across the disk once. That translates into major time savings. All of the reads got done in much less time.

As of today, only drives manufactured by Maxtor and Seagate have this feature. It should not be long for other vendors to follow.

Summary

Choosing the right hard drive for that new computer or to upgrade an existing computer requires some thought, but isn't that difficult when you have a reasonable understanding of the major issues. The table below provides a quick overview of the information in this article

Type	RPMs	Cache Buffer	Cost $/GB	Max Size (GB)
PATA	7200	8 or 16	.50 - .75	400
SATA	7200	8 or 16	.50 - .75	400
SATA	10K	8	2.50 – 3.50	74
SCSI	7200	8 or16	3.00 – 3.50	181
SCSI	10K, 15K	8	2.50 – 3.50	300

WNYLC Web Statistics For March 2005

Total Hits……396,058
Total User Sessions……..50,132
Average Hits/Day
(Monday – Friday)……16,069
Average user Sessions/Weekday……1,893
Number of Pages Viewed……135,729
Average Number of Pages
Viewed Per Day……4,378
Number of Documents Viewed…….83,580

Accessed Using Internet Explorer…….89%
Accessed Using Netscape……..……6%
Operating Systems Used:
Windows 98……..……..16%
Windows 2000……...……..19%
Windows XP……..……...52%
Windows 95……...……<1%
Windows ME……...……1%
Windows NT……..1%
Macintosh……..…<1%
Linux/Unix……...<1%

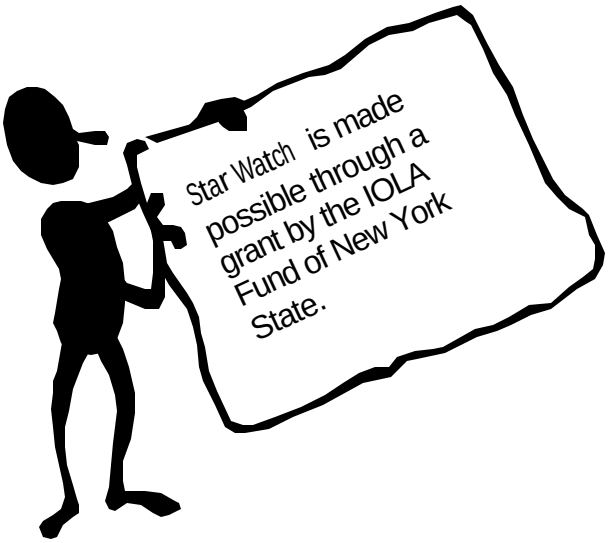

WHO WE ARE

Joe Kelemen - Attorney
Kathleen Lynch - Attorney
Tom Karkau - Programmer
Carly Bouchard - Administrative Assistant

Wnylc@wnylc.com

716-855-0203

www.wnylc.net

Western New York Law Center, Inc.
237 Main Street, Suite 1030
Buffalo, New York 14203