WHY HARDWARE STRIPING

INTRODUCTION

Striping (RAID 0) is the process of dividing the data to be recorded into two or more segments and recording them concurrently to two or more disks drives.
It is done for two primary reasons;

1. To gain throughput performance.
2. To gain volume size.

Of the two, the first is the most prevalent, especially for storing or handling of graphics or editing video information and for real time data acquisition. The second is most common where very large files are being handled in a sequential fashion such as in video recording.

There are two basic forms of striping; by byte where a data block is divided into two blocks of even and odd bytes then written onto two disks, or by chunk where the data is divided into parts, each part being alternately written onto the disks. Of course, any disk I/O of less than a single chunk size is written to only one drive, which then has no performance advantage over a single disk.

Of the two, methods, the chunk striping method is, for most purposes, more efficient and far more commonly used. The end result will or should be data throughput that is near twice that of a single drive while overall capacity is equal to the sum of the individual drives. This is especially true where the chunk size is equal to the hardware block size, generally a characteristic only found in hardware RAID controllers.

Several manufacturers offer hardware striping (RAID 0) controllers. These are most commonly in the form of SCSI to SCSI bridges that attach to the host system SCSI port and connect to the disk drives through one or more separate SCSI ports. The RAID set is presented to the host as a single SCSI drive ID.

Alternatively, several operating systems such as Microsoft NT4, Linux and SGI IRIX have the ability to stripe incorporated into software. For those like Apple Computer's MacIntosh that don't have an OS software striping option, there are independent vendors that offer it.

As you might expect, there are advantages and disadvantages to both methods. Lets examine these:

PERFORMANCE

Software striping can benefit from the use of low-cost standard components. It can also attempt internal optimizations since it knows the OS file structure characteristics. However, this very attribute can be the source of problems whenever the OS or system software is changed. Software striping can, and is often designed to approach the performance of hardware striping.

Hardware striping, in its simplest form, Is optimized for maximum throughput performance and performs its task in this manner, regardless of OS file structure and revision levels. It is insensitive to whatever software is being run, and can be designed such that it will always be able to match or exceed the performance of software striping. There are two basic reasons why this is so.

First is that, unlike software striping, hardware striping does not use host processor resources (overhead). All striping functions (interleaving of sectors and communicating with two or more physical drives, etc.) are accomplished through the use of dedicated processors inside the striping controller and are therefore transparent and impose no additional overhead to the host system.

Second is that software striping requires at least two or more host SCSI transactions (one for each of two or more drives) for every single transaction of the hardware controller. As the limits of the SCSI bus are approached, SCSI overhead for each transaction limits throughput to a greater extent than does the hardware implementation. Some software striping schemes try to overcome this by striping across dual SCSI channels. This increases the cost and complexity of the host system for the additional SCSI channel without really diminishing the overhead.

Let's take a look at a real example of this. A system is required to simultaneously input, display and record 640x480 300kb video information at 30fps in YUV422 format. This requires a sustained data rate to the RAID array of 20 MB/sec. The system uses software striping to a pair of 10,000 RPM Seagate Cheetah drives. This doesn't appear to be a problem since the system has been tested in excess of 28MB/sec of sustained throughput. However, the user encounters excessive frame drops when recording, and when the recorded information is played back, it is fuzzy and incomplete. When the user tries the same test without a simultaneous display, everything is recorded correctly and playback is error free.

Then the user changed to a hardware based RAID 0 (striping) system and repeated the process, the system was able to successfully input, record and display with no errors. Playback was clean with no drops.

In this example, the host processing power was marginal for the required task. It simply ran out of computing power. The next more powerful system was really required to do the job. However, the user

got away with using the lesser system just by implementing hardware striping. The cost for the software vs. hardware striping equation would have included the difference in purchase price of the two host systems.

SCSI BUS CROWDING

Hardware striping units present the striped array as a single SCSI ID, leaving the remaining IDs available to other devices such as backup tapes, CDs, DVDs. Software striping requires that each disk of the array have its own SCSI ID on the host system. Two disks = two SCSI IDs, four disks = four SCSI IDs, etc. Also, not all disks are able to handle the SCSI bus overhead efficiently and may exhibit poor disconnect handling. As a result the SCSI bus performance may be well below the expected speed.

INSTALLATION AND SETUP

installation and setup of a hardware striped array is simple. It is the same as adding a single external disk drive. Set the desired SCSI ID, make sure the cable is terminated properly. THAT'S IT.

Installation of a software striped array starts with determining whether the OS implements striping, and then whether or not it is compatible with the application software and equipment to be used (video boards etc.). After that the decision has to be made between the OS software vs. the various independent suppliers of striping software that will probably result in a more effective solution.

Once that decision is made, then the software must be installed, set up, and then tested to see if it works properly with the system and all of its elements. Each of the drives has to have its ID set and connected to the system and then tested individually. After that, striping is enabled and tested.

REVISION ISOLATION

I sometimes call this the "hidden frustration factor". Take the time to read the notes from the various users groups. There tends to be an undercurrent of concerns, frustrations and problems regarding which disk striping software is going to be compatible with a particular OS revision or a pending OS update (Apple OS 8.5, Microsoft NT 5, etc.) or the latest version of Premiere 5.0 from Adobe or Speed Razor 4.0 RT from in:sync, or the latest Media100 xr system or ReelTime NITRO from Pinnacle. And the list goes on. The solutions being offered always start with checking the web page of the software striping supplier to make sure that the latest version is currently installed.

Hardware striping is immune to this problem. The array is seen by the system as a single, generic SCSI disk drive. The only question that has to be asked is will the system run a SCSI drive?

COST

There is no argument. When it comes to price, software striping wins hands down, even if its third party software. Packages run from 0$ (OS included) to seldom more than $200, whereas hardware striping solutions tend to run in the $500-$1000 range.

But, what is the real cost?

Software striping always adds system overhead, whether it's an incorporated OS feature, or third party software (this includes, of course, the amount of memory used for disk caching and the CPU required for the additional processing). Hardware striping, on the other hand, uses no system overhead. To this extent, the real cost equation has to include the cost of this overhead, and is directly related to the degree that sufficient overhead is available for striping without coming at the expense of the task to be performed.

THE BOTTM LINE

The initial price of software striping is cheaper than hardware, but the real cost has to be determined by the user. Only he can determine the value of the easily installed hardware based system compared to the complexity of setting up and testing the cheaper software solution.

Then, what happens when it's time to update the OS or add a new editing board or production software package. He knows how much his production time is worth. Is the day or days of down time needed to download and install the most recent version of his striping software from the suppliers web site (if it's even available) worth his initial price savings ($500-$1000).? Remember, this isn't even a consideration with a hardware based system.

He spent $6000 to $10,000 for the system and software? He's about to spend $3000-$8000 for a RAID array. Does he really want to have less than optimum performance along with the inherent frustrations in order to save $500-$800?

Is it worth while to convert his present array for a one time cost and not have to deal with the performance limitations and future "hidden frustration factors" any more?