

THE ACADEMY OF PLUMBING 3 – SO YOU WANT TO BE A DATA CENTRE?

Got a digital camera, have you? A big one? With lots and lots of pixels? Did its manufacturer neglect to tell you just how many frames you were likely to shoot with it, how quickly they will mount up, how much hard drive space they take up and how rapidly you will run out of said hard drive space?

What can we store our digital picture files on? Floppy disks? Compact Flash cards? Memory sticks? Nah, get serious. CD-Rs, DVD-Rs and hard drives are the only reasonable options. Of these options, CDs are cheap (20p a throw, these days) but small and not of much use to the average 10+ megapixel shooter. They also require a good deal of operator time to burn, archive, and remove from archive when needed. This operator time has to be paid for, too, one way or another. Really good ones, kept well, can last up to 20 years. Cheap, no-name crap can delaminate and become unreadable within a year. Writeable DVDs can be had for little more than CDs these days and store 6x the amount of data a CD can, but write more slowly and only last a few years before becoming unreadable. They also demand a similar level of operator overhead. So, hard drives are increasingly used because they're fast, on-line, operator overhead is low and when you want to move your data from them to the latest and greatest storage medium, all you have to do is plug 'em in, start copying, and go away for the weekend. The trouble is, hard drives also fail.

A chap called James Wiebe, proprietor of Wiebetech Inc., manufacturer of hard drive enclosures, has written an excellent, readable White Paper concerning the failure rates of different hard drive configurations (<http://snipurl.com/s6mv>) and concluded that individual hard drives have an uncomfortably high rate of failure. What's needed is protection in numbers, and this is provided by RAID (Redundant Array of Independent Disks). A RAID array of hard drives can be configured a number of different ways, and here are those of interest to us.

THE RAIDING PARTY

RAID 0 uses software or hardware to make two or more hard drives appear to us to be one big hard drive. This is how your Big Disk/Bigger Disk/Biggest Disk works. The advantage of this setup is that you get massive drive space and very high speed, because all the drive mechanisms are in use at once with data being striped between them. The disadvantage is that in contravention of what the name implies, no hard drive in the array is actually redundant: lose one, lose it all. So, the failure of any one hard drive in a RAID 0 array causes total data loss.

RAID 1 mirrors hard drives in pairs so that identical data is simultaneously written to both drives, making them exact clones of each other. Writing data to them is a little slower than writing to a single hard disk, because of the processor overhead; reading data from them is fast because most RAID software reads from both disks simultaneously. Wiebe says that RAID 1 offers the greatest data security of any drive configuration. Unfortunately, it does so by using half of your hard drives as backups, so it's not efficient in its use of hardware. It is pretty easy to set up, though, so those of you with dozens of LaCies lined up in orderly ranks across your desks might consider using something like SoftRAID (<http://www.softraid.com/>) to pair them up as RAID 1 volumes for your data archive. The poor hardware efficiency (and 'orrible tangle of wires behind rows of LaCies) has caused manufacturers to create other RAID configurations that make better use of hardware, and the two of most interest to us are RAID 3 and RAID 5.



RAID 3 takes several drives (five in the G-Tech RAID Pro, described below) and treats four of them as a RAID 0 array, striping data across them all. It uses the fifth drive for parity data – essentially metadata about the real data stored on the other drives. Because of this, if any one data drive fails, the data it held can be reconstructed on a replacement drive from the data and parity metadata stored on the other drives in the array. Furthermore, reads and writes are very fast; performance is not impaired if a drive fails (you may not even notice a drive has failed); drives can be hot-swapped, and the whole array can be rebuilt in the background while still being productively used. RAID 3 makes efficient storage usage of 80% of your hard drive space, too, and its data security is comparable with RAID 1. Not bad, eh? The downside is that until recently, RAID 3 devices have been pretty dear, not least because the hardware RAID controller is difficult

RAID resolution

If you're desperate to know more there's a good, succinct explanation of the different RAID types, with diagrams, at (<http://www.acnc.com/raid.html>), or, if you prefer Wikipedia, here (<http://snipurl.com/s6w3>). All four types of configuration are available either for direct connection to a host computer, or as standalone Network Attached Storage devices that simply need plugging into a network and are immediately accessible by other computers on that network. Nothing's perfect, though: all of the formats have their strengths and weaknesses.



to design. So, to try to make things a little more affordable, manufacturers came up with RAID 5, which is similar to RAID 3 except that all drives are used for both data and parity metadata.

In the Buffalo TeraStation Pro (more below), four hard drives can be configured as a RAID 5 array. Again, because of the way data is written, the failure of any single drive mechanism is not fatal and the whole array can be restored. This approach is also rather cheaper than RAID 3 and still manages to make use of around 70% of the hard drive space for data storage. The downside is that RAID 5 is not as fast as RAID 3, especially when being written to (it writes at half the speed it reads) and the array has to be taken out of use while it is rebuilt. It's secure, though: as secure as RAID 3.

COME ON, TELL US ABOUT THE GEAR

OK. There are many manufacturers of boxes like these including Wiebetech, FireWire Depot, LaCie and others, but leaving aside RAID 0 and RAID 1 arrays made in software from loads of LaCies, in my opinion these two boxes represent what are probably the most cost-effective choices currently available.

First of all, the Buffalo TeraStation Pro is a Network Attached Storage device, or NAS, and available in capacities from 0.6TB to 2TB. It's basically a little Linux server and four drives in a box. It does all of its communication via Gigabit Ethernet and is controlled via a

browser-based user interface. In its Pro incarnation its industrial design appears to owe something to the Parthenon, it has swappable hard drives and an LCD to inform you of various interesting tidbits such as its IP address. Its previous non-Pro incarnation, such as the one I own, has the drives mounted within its chassis and an industrial design best hidden in a closet. Nonetheless, its 1TB version offers around 700Gb of RAID 5 storage at around 80p per Gb – cheap as chips for its range of features which include full Windows and Apple file sharing; Users and Group access privileges; built-in schedulable start-up, shutdown and back-up to other TeraStations; Jumbo Ethernet Frames for faster data throughput; USB2 ports for attaching extra drives; compatibility with uninterruptible power supplies; and so on. The downside is that it is slow. The best transfer speed I've seen with my non-Pro is 5Mb/sec write and 10Mb/sec read. Its drive volume is also in a Unix format that can cause problems with Macintosh filenames and file permissions. Also, should it suffer a software problem, the unix drive format and fact that it's a closed system means that resurrecting the data, or even running disk fixit utilities, is somewhat problematic. Nonetheless, it makes a decent medium for a picture archive. Learn more about it here (<http://snipurl.com/s6tl>); learn how to hack it here (<http://snipurl.com/s6tm>) and, if you wish to lend your patronage to an AOP member, get it here (<http://snipurl.com/s6to>).

The G-Tech RAID Pro is also available in capacities from 0.6TB to 2TB (with 3TB coming along soon) and also offers swappable drives, but in every other respect is a different beast. Its drives are hot swappable, meaning you don't have to turn the thing off to swap them out. Its fan and power supply units can also be replaced with ease. It connects to a host computer via FireWire 800 and, being RAID 3, is capable of saturating the FireWire bus with reads and writes of up to 50Mb/sec. And that's about it: it's just a big, fast, Mac (or Windows) -formatted secure hard drive. If you want to share it over a network, turn on Personal File Sharing in your Sharing System Preference and get something like the freeware SharePoints (<http://snipurl.com/s6ub>) to set it up as a network share. In this guise, when attached to a G4 or G5, it reads and writes over a Gigabit Ethernet network at up to 20Mb/sec, much faster than the TeraStation, but of course requires a host computer to do it. Your fast retouching computer can do this without you even noticing and you can even make the G-Tech your Photoshop scratch disk for improved retouching performance. At £1,100 for the 1TB version it works out at £1.10 per GB of storage. You get what you pay for. Find out more here (<http://snipurl.com/s6uq>) and, if you feel so inclined, get it from that same AOP member here (<http://snipurl.com/s6uu>).

Of course, **neither of these boxes will protect you against software failure: if the disk directory corrupts, your data is hosed**, so it's best not to put all of your eggs in one basket. At least have a back-up. The trouble is, a full back-up is the equivalent of RAID 1, tying up half of your hard drive space with back-up data. You can keep this off-site, though, protecting you against fire, flood, burglary and Buncefield. As I said, nothing's perfect... and more about organising your data, including back-ups, next month.

Alright Guv? Those of you still hungry for information can keep up by periodically checking my blog at <http://www.thedigitalplumber.co.uk> where you'll also find my contact details. Where's the Mole, grip?