

THE ACADEMY OF PLUMBING 6 – BACK-UP REVISITED

I've found that many of you are still confused by the plethora of choices when it comes to storage and back-up, so this month I thought I'd recap and expand upon the examples I described in The Academy of Plumbing 4. Furthermore, a couple of new developments – the storage possibilities available in the new Mac Pros and the TimeMachine back-up system to be included with OS X 10.5 Leopard – have changed the landscape somewhat.

THE MAC PRO

Apple's replacement for the G5 has benefited from its cooler-running Intel Xeon processors by having a redesigned case interior including much more room for PCI cards and storage. The Mac Pro can house two optical drives and four internal SATA hard drives, but the interesting bit is in the way the drives are mounted. Unlike conventional hard drive mounts, in the Mac Pro, drive mechanisms are attached to drive carriers which then slide down rails in the chassis and simply click into place: no need to attach cables.

This method of mounting is very similar to the way removable drives are mounted in the G-Tech RAID Pro and open up the possibility of using the Mac Pro itself as your data archive, something I wouldn't have previously recommended because of the difficulty of removing and replacing drives.

Each Mac Pro is delivered with four drive carriers. It's possible to buy the machine built-to-order with up to four Apple 500Gb drives on-board, but to me the Seagate 750Gb mechanisms and, available early next year, Hitachi 1Tb mechanisms (yes, that's a single 1Tb hard drive mechanism) are more interesting. Bought with one 250Gb drive for the boot volume, then reconfigured with four Hitachis arranged as two RAID 1 pairs, the machine would offer two terabytes of secure RAID 1 storage. Buy two extra Hitachis and drive carriers and, at the end of the working day, swap each half of a RAID pair for the spare drive and you have an instant, portable offsite back-up to take home in your pocket. On restart, the RAID software rebuilds the array in the background, keeping everything up to date. At the end of the following day, swap the drives over again. This offers the best of all worlds: onboard RAID protection against drive failure and lightweight, portable offsite back-up. What do you do with the 250Gb drive the Mac Pro came with? Whack it into an external FireWire box and use it for the daily bootable back-up of your boot volume, or whatever else takes your fancy.

What happens when your two terabytes fill up? Buy some more Seagates or Hitachis, some spare drive carriers and put the full drives into storage. It could get even more interesting than that: Apple has a history of generating third-party markets for accessories for its products (just look at the range of accessories available for the iPod from third-parties) so I wouldn't be at all surprised if some enterprising manufacturer produces a G-Tech RAID Pro look-alike able to accept Apple Mac Pro-style drive carriers. This would introduce the possibility of keeping your entire archive on-line, should you wish.

TimeMachine, meanwhile, is going to change the way we do back-up as it operates in the background. If you need to restore deleted files, for example, TimeMachine will enable you to do it simply on a per-application basis, rather than having to sift through back-up disks as we do at the moment. This really is back-up made dead simple. It's early days for TimeMachine, though, because OS X 10.5 Leopard isn't due for release until next Spring and plenty can change in the meantime, but we can look forward to much simpler and, I'd hope, more reliable storage and back-up systems next year.



SYSTEM EXAMPLES

1 Mac Pro-based

- A Mac Pro configured with four 500Gb, Seagate 750 Gb (or, next year, Hitachi 1Tb) internal drives formatted using SoftRAID 3.6 (<http://www.softraid.com>) or later (which requires OS X 10.4.8 or later) as two RAID 1 arrays. Use SoftRAID rather than Apple Disk Utility. There are many good reasons why. Look here: (<http://tinyurl.com/e6joo>). Choose one array for the boot volume and partition it into two, creating a 100Gb boot volume with the remainder as work area. The second array becomes your primary data store and Photoshop scratch volume.

- A pair of extra drives mounted in Apple drive carriers to use as the daily or weekly offsite backup.
- More drives and carriers to replace your primary storage array when it fills up.
- A housing to put your spare drives in, when such a thing appears on the market.

The only downside to this approach is that your drives will only work in a Mac Pro: they can't be attached to anything else unless you remove them from their carriers and put them in FireWire cases. However, when my speculative third-party drive housing shows up on the market, this problem gets solved. It's a very neat and tidy solution to most of the practical problems associated with devising a storage system. What do you do with that collection of 250Gb LaCies cluttering up your desk? Use them as off-line RAW file storage or, if you're strapped for cash and the drive mechanisms are SATA, take the drives out and mount them inside the Mac Pro instead. You could also donate them to the AOP who would be more than happy to put them to use.

2 Non-Mac Pro-based

- A large RAID 3 or RAID 5 box such as a 1 or 2Tb G-Tech RAID Pro (<http://snipurl.com/s6uu>), a United Digital FireStar (<http://tinyurl.com/htap6>) or similar. If one of these isn't enough, get two. The G-Tech is currently in short supply so you may have to look around and wait a bit to get one.

- Another cheaper drive of equivalent capacity to act as back-up for the RAID box. This could be a LaCie Big Disk Extreme (<http://tinyurl.com/eklx4>), Bigger Disk Extreme (<http://tinyurl.com/gb5j8>) or similar.

- Ideally, another one to keep off-site. Swap them over weekly to keep them up-to-date. These drives are not too bulky or heavy to put in a bag and take home on the Tube. I'd do it.

- Recycle the external drives you currently have as off-line RAW storage and bootable back-ups for your Macs. This is good practice to prevent your relatively expensive, valuable RAID box filling up too quickly.

This setup will remain useful when you eventually move on from your G4 or G5 to a Mac Pro.

3 Budget

- For your main data archive use a pair of LaCie Big Disk Extremes (<http://tinyurl.com/eklx4>), Bigger Disk Extremes (<http://tinyurl.com/gb5j8>) or similar on a twice-daily back-up. Recycle the external drives you currently have as off-line RAW storage and bootable back-ups for your Macs. Get a third big drive as soon as you can afford to for off-site backup.

Big RAID 0 drives such as the LaCie Big/Bigger Disk Extremes are much less reliable than RAID 1, RAID 3 and RAID 5 devices, which is why it's important to have an up-to-date back-up should one fail. Use ChronoSync (<http://snipurl.com/pn05>) to do your back-ups and get in the habit of manually syncing the drives every time you add files to your main data store.

4 Scraping the barrel

- Get a copy of ChronoSync, buy pairs of 250Gb or 300GB drives on-line for less than £100 each, define one as a data store, the other as its back-up and run back-ups daily and whenever you add files to the data store. Daisy-chain them using FireWire 400 to avoid the dreaded FireWire 800 corruption bug. Get a small FireWire hub if need be.

Even this approach isn't much cheaper than Budget, above, because by the time you pass the 750Gb point it would have been cheaper to get a pair of Big Disk Extremes in the first place which you would have been able to run at twice the speed using FireWire 800.

WHAT, NO TERASTATION?

That's right: there's no Buffalo TeraStation in the above set-ups. That's not because the TeraStation isn't a good box and I still think it to be a good idea as a location for your picture archive, but in practice most users I know of are irritated by its slow transfer speeds. It can also be problematic if it suffers a software failure because you can't use your Macintosh (or Windows) disk utilities on it. Nonetheless, mine soldiers on without missing a beat.

IN SUMMARY

Back-up is important. If you've not been bitten by data loss yet, you will be at some point. Obsolescence-resistant storage systems that allow you to do your work and get at your image archive efficiently and with a minimum of fuss are also important. Choosing and setting up these systems can seem daunting at first

And there's more

AOP APPEAL FOR REDUNDANT MACINTOSHES

The appeal is still running until the 15 of October, folks, so let's be having 'em. Our long-suffering staff are already sending smoke signals instead of e-mails and with autumn coming on it's going to get windy.

A NASTY OS X FINDER BUG

I've now had two clients lose data as a result of a nasty bug in the OS X Finder. I daresay Apple would claim it isn't a bug and that the Finder is working as specified, but in that case as far as I'm concerned it's a serious design fault. It's to do with changing your mind when using the Finder to burn CDs or DVDs. Here's what happens:

- You have a folder of files to burn sitting on your Desktop. You insert a blank CD and drag the files into it. You think (and the whole process encourages you to think) that you've actually dragged your files into the blank CD but in fact you haven't; the Finder has made aliases to your files.
- You then decide not to burn the CD or DVD. There could be a good reason for this: you may have dragged in too many files to fit. So, you drag your 'files' back out to their folder again.
- The Finder asks you whether you want to overwrite the existing files. Not really understanding, you say OK. The Finder then overwrites your files with aliases of those files.
- That's it: they're gone, because Undo (Apple-Z) doesn't work in this instance.

What you actually do to get out of burning the CD or DVD is to select all the aliases in it, drag them to the trash, and then drag the disk icon to the Trash. Intuitive? No, I thought not. Why is this a bug or design fault? It's because the aliases created by the Finder have identical names to the files they refer to and don't have 'alias' appended to their name, coupled with the fact that OS X allows you to overwrite one kind of object (a file, folder or application) with a completely different kind of object (an alias), just on the basis of it having an identical name. This is clearly nonsense. Even if there were a good, traditional UNIX reason for allowing such overwriting (which I doubt), it obviously never makes any sense or is ever useful to allow original data to be overwritten by an alias pointing to it, because what are you left with? A useless, orphan alias. OS 9 didn't do this: you couldn't overwrite a file with a folder and vice-versa. Yes, it can be argued that you've been warned and agreed to the overwrite operation but the system should be more user-friendly than this and not mislead you in the first place. Apple needs to get this fixed.

but, like so many things, are straightforward in retrospect. If you don't feel up to the task of selecting and installing this stuff yourself, **get someone competent to help you.** Ahem...

Running SMARTReporter are you, sir? No? Oooh, it's going to cost you... 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.