

THE ACADEMY OF PLUMBING 4 – BACK-UP... THIS ONE'S SERIOUS

Right, it's time to cut the hilarity and get down to the nitty gritty. It's a truism that there are only two kinds of computer user: those who have already lost data, and those who are about to. If you don't want to be among the latter, read on and act on these words, or face a >£650 bill for data recovery. Got your attention? Good.

The primary, over-riding characteristic of digital data is its ability to produce identical copies of itself. Not just 'very, very similar', which is the best the analogue world can offer (analogue – analogous: comparable, parallel, similar, like, akin, corresponding, related, kindred, equivalent), but identical. This wonderful mixed blessing allows us to create genuine, identical copies of an original whilst simultaneously devaluing that original itself. The second most important characteristic of digital data is that it is entirely dependent upon electro-mechanical gubbins to house it, and should said gubbins go south, sayonara to the data, as my wife might say. So, looking on the bright side, how do we go about setting up a storage system that allows us to keep our picture data organised and easily available to us, preserve it if the storage devices fail, and backed up just in case they do?

NOTHING'S PERFECT

Please bear this in mind, because all storage and back-up systems have their drawbacks and it's going to be down to you to choose the system that most suits your working style and personal habits. There's no point in choosing something like tape back-up if you can't maintain the habit of routinely changing the tapes and an unused or improperly-used back-up system is worse than useless. Let's start with some terminology.

- A mirror consists of two or more drives whose contents are maintained in an identical state at all times, because identical data is always written simultaneously to all drives in the mirrored array. If you read last month's column you'll know that RAID 1 mirrors two or more drives together and provides the very highest level of protection against hardware failure, because should one drive fail, its contents are mirrored on at least one other drive. Some RAID 1 mirrors, such as those created by Apple RAID (part of Apple Disk Utility) and SoftRAID (<http://www.softraid.com>) can keep going without missing a beat if a drive fails.

So, the advantage of mirrors is that you lose no data at all if a drive fails, but the problem with mirrors is that they offer no protection at all against software failure. If you inadvertently trash some files or corrupt the disk directory, all drives in the array are equally affected and your data gone.

- A backup consists of two or more drives whose contents might be identical, but are usually similar but not identical. This is because back-up software copies data from a master drive to a back-up drive (or tape, CD, DVD, NAS) at intervals, leading to a time lag between the current state of the master and its state when the last back-up was carried out. Depending upon how much work you're doing and how often you back up, the time lag can be quite small and the discrepancy between the master and back-up kept to a minimum.

The advantage of a backup is that it protects against hardware failure (your data exists in at least two places) and software failure (the time lag means that inadvertent file erasure or corruption of the master won't affect the back-up until the next back-up operation, giving you a chance to recover the situation). The disadvantage is that should your master fail completely, you lose

everything you've done since your last back-up was made. This isn't fun if your last back-up is six months' old.

You can see a sliding scale here: at one extreme, mirrors offer complete hardware but no software protection; weekly back-ups offer protection against both and give you plenty of time to realise that there's been a software problem, at the cost of losing everything you've done since the back-up. Most back-up software allows you to create frequent back-up schedules, even down to a few minutes between back-ups, so where you set your back-up frequency depends upon how much data you can afford to lose and how quickly you think you can spot data loss or corruption.

WHAT SHOULD I BACK UP?

In *The DAM Book*, Peter Krogh reckons you should keep at least three back-ups on two different media. He recommends two drives and one DVD. Fair comment and who am I to disagree, but you'll have read my opinion concerning DVD back-ups last month. What he also reveals is a general principle concerning data: there's no need to keep data online if you don't require instant access to it. He DVD-archives his RAW files, reasoning that he rarely needs to open them if he can create derivative files from his developed, retouched TIFFs or PSDs, and he has a point. Whatever you do, don't back up more data than you have to. Be ruthless with your selects. If you're storing pictures for clients, find a way to charge 'em for the privilege. I also think it useful to keep one or two daily complete bootable backups of each of your Macintoshes' boot volumes. These allow you to start up another Macintosh as if it were your own and represent cheap but priceless insurance, especially when shooting on location. See the box-out for more. A daily backup is good; servers can probably be backed up as often as hourly. Tethered shoots should have the shoot folder backed up very frequently indeed; downloads from CF card should be made to two drives simultaneously by using Photo Mechanic (<http://www.camerabits.com>) or Image Ingestor (<http://snipurl.com/teei>) and the back-up drive stored until the job is delivered.

BACK-UP HARDWARE

Last month we discussed different types of RAID arrays, their strengths and weaknesses, and looked at a couple of tasty high-capacity boxes. But which should we choose? How should we back them up? Based upon last month's findings I recommend the following:

- Use a high-reliability device for your main data store. This can be a RAID 3 or RAID 5 box such as the Buffalo TeraStation, G-Tech RAID Pro or similar. See last month's column for mini-reviews. It could also be a drive housing from WiebeTech (<http://www.wiebetech.com/>), FireWire Depot (<http://www.fwdepot.com/>) or others, filled with 750Gb Seagates arranged as RAID 1 pairs, or even an Apple Xserve RAID similarly configured.
- If you can't stretch to one of these, use pairs of single high-capacity external drives arrayed in software as RAID 1. They'll be more fiddly and, because of the wires and power supplies, less reliable than a big RAID box but the RAID 1 mirroring will protect you against data loss.
- Whichever you choose, back it up often to another, possibly cheaper device of equivalent capacity. A Big/Bigger/Biggest Disk will do the job.
- Use a second Big/Bigger/Biggest Disk as an off-site backup. It's

UPDATES: Of course, us plumbing types are accustomed to dealing with snagging lists before a job is finished, and this job is no different. So, here are a couple of follow-ups to items in previous columns.

STARTING UP FROM A BACK-UP

A procedure I glibly mentioned in both the May and June columns and again this month, without providing you with an inkling of how to do it. Simple. If, on a FireWire disk, you have a bootable backup of your boot volume (the drive you normally start up from), you just plug it in, wait for it to mount on the desktop, open System Preferences/Startup Disk, select it, and restart. Done. If your internal drive has crashed and you need to start up from your back-up, hold down the Alt key after the restart chime. You'll get a blue screen and your external back-up (if connected) will appear eventually as an icon button. Click to select it, then click on the right-hand arrow button. Done.



There are a few flies in the ointment, though:

- You can't start up an Intel Macintosh from a PowerPC Macintosh bootable back-up, and vice-versa. This is because Intel Macintoshes have a different disk partition scheme to PowerPC Macintoshes. Some heavy-duty Unix geeks have actually figured out a way to make a universal boot volume that will boot both kinds of Mac, but the procedure is not for the faint of heart. This anomaly will probably change with OSX 10.5 Leopard, due before the year's end.
- You can't start up from an un-powered, miniature 'mobile' disk. Most of these don't support booting. To boot up, the disk must usually be a normal-sized, powered type.

relatively easy to take home with you or, if you work from home, to leave with a pal. Update it weekly by rotating it with the one already in your studio. If you live and work somewhere prone to hurricanes, fires, floods or politically unstable, perhaps consider sending cheap, lightweight back-up drives to some friendly colleague elsewhere at regular intervals.

- If you want to do a Peter Krogh and make DVDs too, I won't stop you.
- Think about keeping your RAW files, previous years' accounts, etc. offline on large-capacity single drives. 250Gb LaCie's can currently be had from DABS Direct for around £90 each (<http://snipurl.com/te9>). Leave them on a shelf, but not for too long: drive spindles can seize up if they're not exercised, so connect them up and check them from time to time.
- Alternatively, consider using a tape drive for your offline back-up requirements. Most of them work with Retrospect and can provide very cheap storage per Mb, but are more demanding and fiddly to use than hard drives.

- Some applications will require re-authorising before you can run them. Photoshop CS2 will bleat that it is no longer authorised, won't run and will insist you re-authorise it somehow. If you have an Internet connection and haven't run out of authorisations, fine. If not, a tedious call to Adobe Customer Support will be necessary and is not guaranteed to be successful, something to think about if you rely upon Adobe Camera Raw. The phone number will be on CS2's splash screen. You're a lucky person if you happen to be using a site licence, which doesn't submit you to this nonsense. And nonsense it is because, for example, Phase One use a much more photographer-friendly system of licence control. Whereas Capture One is also likely to require re-authorising, their process is a little more realistic than Adobe's in that if you can't make a connection to re-authorise, Capture One will happily revert to fully-functional 30-day demo mode, meaning you're not caught out. The wise person will have all of their serial numbers, customer support numbers, usernames and passwords in one accessible place. I use and can recommend Web Confidential (<http://snipurl.com/qmde>) for this purpose, but there are plenty of alternatives.

APPLEJACK AND INTEL MACS

The very wonderful Applejack (<http://snipurl.com/qh2o>), a free start-up shell-script for troubleshooting and repairing system problems, doesn't currently run on Intel Macs. So don't try it.

RAID 3 WORKS

Since writing last month's column a G-Tech RAID Pro belonging to a client of mine has suffered a failed drive. The drive suddenly acquired a red light next to it; when the unit was shut down and restarted the light went off and the G-Tech Utility software simply reported 'no drive present' for that specific drive. The RAID Pro continued to function throughout all of this without missing a beat, so it did perform as advertised and the RAID 3 format preserved all of the data on the unit. A 1Tb LaCie Big Disk (internally, a RAID 0) would have permanently lost everything under these circumstances. All of the other G-Tech RAID Pro's I know of are performing perfectly.

FIREWIRE PROBLEMS

FireWire drives can come and go a bit. As a general rule, if you're getting odd behaviour of any kind with a FireWire drive, power everything off, wait a couple of minutes, then boot up again to reset the FireWire bus. FireWire 800 has its own problem: an obscure bug means that if you have more than two FireWire 800 drives daisy-chained on one FireWire bus you are likely to get data corruption. The solution is to add more FireWire 800 buses by adding FireWire 800 PCI cards such as the LaCie card available from the AOP Apple Store.

BACK-UP SOFTWARE

There's loads of it out there. ChronoSync (<http://snipurl.com/pno5>) has a good reputation and offers a wide range of features at a reasonable price, but it can't make bootable back-ups (see the box-out) and I personally have had a couple of problems with it. Synchronise Pro (<http://www.qdea.com/>) has a reputation for being powerful and bulletproof, but is expensive and has a rather onerous licensing model. Retrospect (<http://www.dantz.com/>) used to be king of the hill, but its reputation has slipped of late as it has suffered several changes of ownership and its developers appear to have taken their eyes off the ball.

A good place to search for back-up software, and software in general, is VersionTracker (<http://www.versiontracker.com/>). A VersionTracker list of back-up software is as (<http://snipurl.com/te95>).

Read, and then choose. None of them are perfect; one or more of them will do the job for you. I currently use a combination of ChronoSync, SuperDuper! (<http://www.shirt-pocket.com/>) for bootable back-ups and Déja Vu (<http://snipurl.com/te9c>) for backing up some client systems.