

THE ACADEMY OF PLUMBING 14 – STORAGE REVISITED

As the AOP Digital Season recedes into the mists of time, it has become clear to me that storage and backup remains problematic for many of us. Things have changed somewhat in this area since I last wrote on the subject, so I thought it was time for a reappraisal of what to do, how to do it and what to do with it.

It's a digital truism that data doesn't exist unless it exists in at least two places simultaneously. To this can be added the notion that online data can't be considered safe unless an offline version of that data also exists. To cut the jargon, this means that for belt-and-braces protection against data loss, ideally you want two identical copies of the data connected to your computer (online) with a third copy on a shelf somewhere (offline). How best to achieve this without drowning in external hard disks, their cabling spaghetti, the onerous task of managing all this stuff and draining your bank account? You get a RAID box and fill it with hard disks RAIDed together.

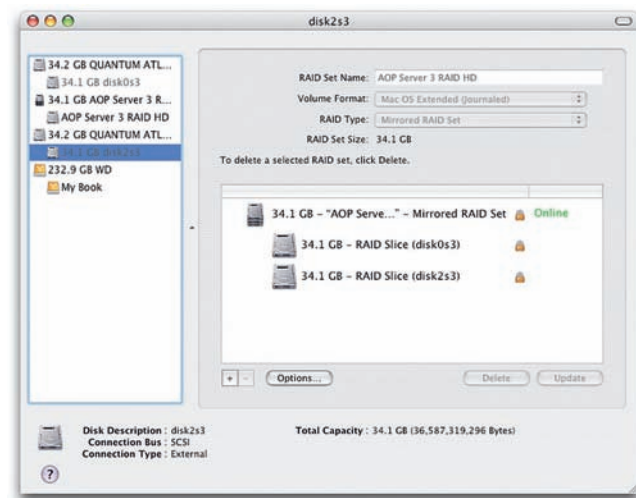
RAID EXPLAINED, AGAIN

Actually it's all here (<http://en.wikipedia.org/wiki/RAID>) but I'll give you a quick and simple précis. RAID stands for Redundant Array of Independent Devices. 'Array' speaks for itself: it refers to several devices arranged in a system to create one larger 'virtual' device. The 'redundant' bit means that at least one of these devices is redundant as far as keeping your data complete is concerned because it either acts as a spare copy, or contains metadata information enabling you to rebuild that data should one of the devices fail. First let's be clear about the differences between devices, partitions and logical volumes. For our purposes, a device is a single physical hard disk mechanism. A volume is what you see mounted on your desktop. A device can have one or more volumes on it, which is what happens if you split your internal hard disk into multiple partitions: they all individually mount on the desktop but live on the same physical hard disk. RAID takes multiple devices and arranges them into one or more logical volumes. There are several types of RAID, of which we're only interested in a few:

- RAID 0, also called striping, is the odd one out in that it offers no redundancy: it takes multiple disks and creates only one volume from them, using their combined capacity. So, if you take two 500Gb disks and array them as RAID 0 you'll get a single 1Tb volume mounted on your desktop. It's very fast, but offers no protection against disk failure because if one disk fails you lose everything, irrecoverably.
- RAID 1, also called mirroring, takes two or more devices and makes one volume from them, which is the size of the smallest device. Take two or three 500Gb disks, RAID 1 mirror them and you'll get one 500Gb volume on your desktop. The point of this is that all data is simultaneously written to all of the disks in the array, so the contents of the disks are always identical. Writes to disk happen at a similar speed to writing to a single disk; using SoftRAID (below), reads are very fast because all disks in the array are read from simultaneously, making RAID 1 as fast to read from as RAID 0. RAID 1 offers simple, complete protection against disk hardware failure at the cost of only using at best 50% of your available hard disk capacity.
- RAID 3, 4 and 5 all try to work around this limitation and offer protection from disk failure (but not RAID controller failure), large volume sizes plus high read speed. They do it by taking

multiple disks, making one or more volumes on them and writing the data to the disks in a clever way such that should one disk fail, its contents can be reconstituted from the other disks in a process that happens in the background, transparently to you. You can get up to 80% of your disk space used for usable storage (as opposed to safety copies) using this format, but at the expense of added complexity and cost compared to RAID 1.

The easiest forms of RAID to set up are 0 and 1, both of which can be done in software using Apple Disk Utility or SoftRAID. The other forms require dedicated hardware controllers that are usually part of the boxes themselves.



STORAGE HARDWARE

Well, the obvious thing to do is to take that row of 16 LaCies strung out along the back of your desk and RAID them together in various ways. Go ahead, but don't expect the system to be terribly reliable because its spaghetti of FireWire cables and power supplies rather militates against that. It's better and cheaper in the long run to get a dedicated multi-disk box of one form or another because it usually has a better build quality, much better power supply components, fan cooling and of course requires far fewer wires to connect it up. Let's look at what's available.

NETWORK ATTACHED STORAGE (NAS) devices are actually little headless file server computers with hard disks in them, and attached to your Ethernet or wireless network rather than directly to your computer. You can get them as single disks or as RAID 5 units housing four or more disks. The Buffalo TeraStation (<http://tinyurl.com/39vlza>) is one such. They're often excellent value for money but can be slow in that they only write and read data at network speeds, and can be troublesome to fix if they go wrong because you can't use your Mac or PC disk tools on them directly. They never use the Mac disk format, so you can run into problems with file names and special file permissions.

LOCALLY-ATTACHED RAID 3, 4 OR 5 boxes such as the G-Tech RAID Pro connect directly to your computer using FireWire, USB or eSATA (an external variety of the connection your internal disks make with your computer's guts) and offer much higher transfer speeds. This is because they are primarily targeted at video editors who require really high-speed storage and delivery of a lot of data: the design goal is to make much bigger, faster volumes available than any single hard disk could offer. Therefore, they tend to be quite expensive but are a good choice if you need to keep a lot of images

reliably available online because you're running a picture library, for example. They offer a lot of hardware redundancy with hot-swappable disks (you don't have to turn anything off to remove and replace their hard disks) and other components such as power supplies and fans being quick and easy to replace. Both the G-Tech RAID Pro (which is no longer available) and Buffalo TeraStation were described in IMAGE 370.

LOCALLY-ATTACHED JBOD BOXES. This is an interesting area and the one that has changed the most since my last column on this subject, primarily because they've become much cheaper and more easily available. JBOD means Just a Bunch Of Disks (nerds, don't you just love 'em?) and that's what they are: a box containing multiple hard disks that you can arrange as you wish. In these boxes, bare hard disks are screwed to disk carriers or 'sleds' and simply slide into slots in the box casing. They make a push-fit connection to sockets inside the back of the box and click into place. Dead simple. It takes about two minutes to screw a disk into a sled; spare sleds are easily available and the disks can simply be clicked in and out while the whole system is running. The boxes themselves attach directly to your computer using the same variety of connections as the RAID 3, 4 and 5 boxes and can be shared over your network just by enabling Apple Fire Sharing in your System Preferences/Sharing preference pane. There are now several manufacturers making these; in my opinion those made by Sonnet Technologies (<http://tinyurl.com/2okfgf>) offer good build quality and excellent value for money with the virtue of being easily available in the UK. Funnily enough, the Mac Pro can also be considered to be a JBOD because of the ease with which one can remove and replace its internal hard disks, although Apple has somewhat muddied the waters by making the disk sleds slightly different in the new 8-Core Mac Pro compared to the others in the line, and 3rd-party housings for these sleds have so far failed to materialise.



Of Sonnet's range, the Fusion D500P and Fusion D400Q enclosures are probably of most interest. The Fusion D500P has space for five disks and connects by eSATA. It features port-multiplication: normally, each SATA hard disk requires its own connector cable, but port-multiplication enables the box to connect all of its disks to a suitable eSATA PCI card in your computer with just one wire. Simpler is definitely better as far as wires are concerned. Sonnet make appropriate cards: the Tempo E4P for PCI Express as found in the Mac Pro and last-generation G5s, the Tempo X4P for PCI-X as found in most G5s and the Tempo SATA ExpressCard 34 card to whack into a MacBook Pro, which enables you to connect up to two Fusion D500Ps to it. Those of you with PowerBooks can get a CardBus card made by Wiebetech (<http://tinyurl.com/39a7eh>) to do the same job.

Sonnet's two PCI cards each have four connectors; it's feasible to have four Fusion D500Ps connected to one card, each box being filled with five 1Tb disks. That's 20Tb of hard disk storage available online at once, which ought to satisfy you for the next year or two. The Fusion D400Q offers four disk slots but doesn't require any host cards



SATA-On-the-Go
TEMPO SATA EXPRESS 34

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.

Recover your stolen mac!

I recently installed a nice new iMac and MacBook Pro for a client in the countryside. Less than a week later, the naughty boys with the white van kicked down their front door and nicked 'em. Unfortunately, neither had Undercover (<http://www.orbicule.com/>) installed. Once activated, this rather clever software not only 'phones home' (that is, Orbicule's servers) with the stolen Mac's IP addresses as soon as it establishes a network connection, it also sends back screen grabs and video snapshots of the phyzzog of whomever might be leering at it, ie the thief, fence or hapless eBay purchaser. After a bit it will start faking hardware failure and prompt you to take the Mac to a dealer to get fixed, and even go so far as to shout 'I'm stolen!' at you. Very good, and if those two Macs had had it installed I'm sure they'd be back with their rightful owner by now and avoided the unpleasant insurance wrangle that ensued. I've learned my lesson and installed it. The publisher is a Belgian company, too, which makes a change from the US dominance of innovative software. Maybe they learned a thing or two from their countryman Inspector Poirot.

as it has all four sensible connection formats on-board, making it attractively versatile. The Fusion D500P is available empty for around £300; the D400Q is around £370; the PCI cards are around £160 each and the laptop cards are about £60 - £70. Spare disk sleds are around £20 each. Considering that it's now possible to get decent 500Gb high-speed eSATA II disks online for less than 60 quid before VAT, it now becomes plain that as soon as you start looking at having proper backups, especially offsite backups, these boxes make sense from the economic, reliability and ease-of-use points of view.

RAID AND BACKUP SOFTWARE

All right, then. We've decided to stump up for one of these low-cost JBOD boxes. Let's say it's a Fusion D500P. How do we use it? Like this.

- 1) Load it with five disks.
- 2) Use SoftRAID (<http://www.softraid.com>) or Apple Disk Utility to either
 - a) create one RAID 1 trio and leave two spare disks; or
 - b) two RAID 1 pairs and leave one spare disk.

If you've decided on 2a, the three disks in the RAID 1 array give you complete hardware protection plus a third disk to pop out and keep offsite. One of the non-RAID spares acts as a time-delayed backup to guard against software failure: set up ChronoSync (<http://tinyurl.com/36yy9>) to make a daily backup for you. The other is set up to be a bootable backup of your internal boot disk, using either SuperDuper! (<http://www.shirt-pocket.com>) or Carbon Copy Cloner (<http://www.bombich.com>). If you've decided on 2b, the spare disk acts as a ChronoSync backup to one of the RAID 1 volumes. Recycle an external disk as the backup for the other.

3) When all of these volumes fill up, remove them from the box, split the RAID 1 pairs or trios, keep one in a box on a shelf and another offsite. Recycle the backup disk to become the backup for the next pair or trio of disks you install. Instead of a shelf of DVDs on your studio wall you'll end up with a shelf of hard disks, each sealed in a plastic freezer bag along with a little sachet of silica gel, labelled and stored in a VHS video box. Alternatively, define a pair or trio of disks as a working volume for a particular client, or shoot, and remove and replace it wholesale whenever you need to. The box starts to work like a big cassette recorder.

Sounds good?