

## Digital Image Storage/Archive

Copyright Brian Spangler - All rights reserved.

Like many of you I have a hard-drive of folders with each containing a single 'Master Image' that is used for all customer printing, submissions, transfers to film and web-based creations. Given today's requirements for True Color Imaging, 48 bits per pixel, 5000dpi and so on...it is very easy for a particular Master Image Folder to occupy from 100 to 150 megabytes of hard-drive space.

At that size it only takes a season of photography before even the most relaxed of us start to become anxious about the amount of hard-drive space left on our machines. This has led me to begin investigating a more appropriate and more commercial level of digital image archive management.

In the data management industry the general approach to information storage is to incorporate multiple levels of redundancy into the overall data workflow so that at no time does there exist a single point of failure that will cause a loss of valuable data. One NPN member I spoke with accomplishes their storage redundancy by utilizing multiple hard-drives and doing full or incremental backups. All of his hard-drives contain the same version of each image and each hard-drive remains synchronized as new digital images are added. The benefit to this approach is that if any one of the hard-drives fails, you are not compromising the safety of your work; you simply replace the drive and re-synchronize the image data to the master drive containing the original images.

At a commercial level you can purchase a RAID Data Storage System. RAID, Redundant Array of Inexpensive Disks, is a method where the image data is spread across several hard-drives thru techniques called disk striping and disk mirroring. This achieves both redundancy as well as allowing the easy addition of new hard-drives to increase your total storage capacity. To the computer this array of disks appears as one large volume or hard-drive and functions no differently than any other drive in your system. However, a key difference is that if any single RAID component fails your PC will continue running and working with what it sees as just another drive. That is the beauty of RAID -- by utilizing many smaller disks and building redundancy the RAID system can present itself as one large volume while ensuring continued access even when a part of it has failed. Another valuable benefit is RAID's ability to add additional storage devices on the fly. If your logical volume is filling and there exists physical room in the hard-drive bay, then a new hard-drive is all the doctor ordered. Once you pop the new hard-drive in the available drive bay, the RAID system will assimilate the new guy and soon the drive is working as part of the large volume.

For the owners of this mega hard-drive we are afforded a hot-swappable hard disk array with built in redundancy to allow swaps of our drives when we want to add more storage or when an individual drive has gone bad. When failure of a drive does occur it can be pulled from the array and replaced without ever interrupting ongoing image workflows.

Setting up a RAID system today is as easy as a trip to the store or a visit to the website of Dell, Gateway, etc. The types I am considering, and would recommend, are those that are external, self contained and have their own power supply, such as the Accusys 7500 IDE RAID System. The 7500 was found online for as low as 249.00.

Although you can build your own by purchasing a RAID Controller and installing multiple drives in your machine, this will only add overhead to your system that may or may not already be taxed by the work you are doing in Photoshop, Paint Shop Pro or equivalent. Leave the PC to manage your image editing and workflows and use the separate RAID box as an external device to store your images safely and reliably.

The methods of 'redundant hard-drives' and RAID described so far will help safeguard your images, but a greater threat to your data is location, location and location. When your time-invested work and creativity resides in one location there is no amount of redundancy available that can save your digital images from theft, fire, flood or even a malicious power surge from a springtime lightning strike -- in an instant it can all be gone. This potentially devastating scenario exposes the critical need to store backups of your works at an off-site location.

A relatively new backup solution, which has become consumer-affordable, is the DVD writer. For both the PC and Mac, good quality DVD mastering hardware is now available for under \$500. But, before you run down to your local Best Buy to pick one up, let's do a few calculations to see just how feasible DVDs are for our needs. With the redundant hard-drive approach you do not need to back

up any of the drives used for redundancy. Remember that the additional hard-drives are used to back each other up so all the hard-drives contain the same data. For storage to DVD, let's say each of your three drives is a 256 Gigabyte unit, for a total capacity of 768 Gigabytes, your goal is to only backup the primary 256G hard-drive. Since each DVD holds roughly 5G, you would be looking at 256/5, or about 52 DVDs. If this approach is doable then I would consider writing the DVDs every couple of weeks and placing them in a safety deposit box or ship them to a friend in another state who is willing to store them.

I should mention that today there is software available that provides you with the ability to exercise your DVD writer and perform incremental backups. Whether or not this is helpful depends on the volume of data you will be backing up.

If you are just starting out and your current data volumes do not exceed the size of one or two DVDs then this is a good solution to begin with. The only way to make this work is to stay in touch with the amount of new data you are adding regularly and to perform a backup when the total size of your new image data will fit a DVD or two. Your labor will be limited to the backup of just one or two DVDs periodically which is a minimal effort for achieving redundancy and off-site storage.

Although this is a solution, in my opinion it is not practical if you already have a large amount of data already stored. Considering the 256 Gigabyte drives, a typical burn time for a single DVD is about 15 minutes, so at best you will need to plan 12 to 15 hours to backup your 256 Gigabytes! Not to mention that the cost for the DVD disks will come in at between \$150 to 200 dollars. Granted, the above scenario outlines a photographer with image storage requirements at 256 Gigabytes of data. Is that typical? Well, although I am not yet at that level, when I consider each of my Master Image Folders, containing files .txt, .psd, crw, .xmp, etc., and requiring anywhere from 20 to 100MB, it is not inconceivable that in a worst-case scenario (100MB each) my 256G will fill up at roughly 2,560 Master Image Folders. I am sure there are many of you already exceeding this number in a given season. So, what are some other alternatives? One possibility would be off-site archiving via a web-based service that can link to your machine and copy the data -- all appropriately secured and encrypted.

Off-Site archiving can be one of the easiest solutions to use since you do nothing but let the machines talk and wait until they have completed the transfer of data. This type of service does require a good broadband connection that has a high upload speed. The cost, however, is significant, as sites such as [www.globaldatavault.com](http://www.globaldatavault.com) have generally intended on servicing business data which is far less sizable than a digital photographer's image data. On average these services cost about \$300.00 per 30 Gigabyte each month.

Another site's service, [www.verio.com/powerplatform/services/storage.cfm](http://www.verio.com/powerplatform/services/storage.cfm), was established to offer both full and incremental backups of data volumes ranging from 35 Gigabyte up to 1 Terabyte. I have forwarded an email requesting pricing information but have received no reply.

iBackup ([www.ibackup.com](http://www.ibackup.com)), is the last site I looked at and in my opinion they offer one of the slickest web-based services for storage of your data. At \$800.00 per 100G a month, this cost could be an issue for many of us.

With the alternatives outlined and their costs listed you may or may not have enough information to make a decision. For now I am leaning in the direction of establishing a RAID system with hot-swappable drives. Why you ask? Simple, hard-drives are quite cheap...each time I want to store new data off site I will purchase an external drive (one that can accommodate the volume size of the RAID system). The external drive is plugged in to the computer via the Firewire Port or USB version 2.0 access port, once the external drive is visible from my desktop I can start backing up a single instance of all my work from the RAID system. Keep in mind that to your Mac or Windows Desktop, the RAID's array of hard-drives appears as one big logical drive so the backing up is just a copy from one volume to another.

Once the copying between drive volumes is completed I will have in my hands a new backup that I will then carefully pack away and store in a safety deposit box at my bank. When it comes time to do another off-site backup I will purchase a second external drive, repeat the backup process and replace the external drive in the bank's safety deposit box with the newer backup. In this process I will eventually need to increase the volume sizes of my two external hard-drives to match any volume increases in my RAID system.

I hope you find this article informative and I would ask that you share your storage setups with the rest of us here at NPN so that we can all learn a little more today.

*Editor's note* - Be sure to visit Brian's web site at [www.outdoorgrace.com](http://www.outdoorgrace.com).

Brian Spangler - NPN 723