



## TMUG Officers

**NOTE: Some officer e-mail addresses will be changing soon!**

### President

Duane Reaugh  
pres@tmug.org

### 1st Vice President

Allen Emory  
vp1@tmug.org

### 2nd Vice President

Melanie Crain  
vp2@tmug.org

### Secretary

Ruth Landa  
secretary@tmug.org

### Treasurer

Stefan Jeglinski  
treasurer@tmug.org

### Web Master

Frank Crigler

### Postmaster

Paul Lemieux  
plemieux@tmug.org

### Get Info Editor

Lisa Colwell  
editors@tmug.org

### Listserv Administrator

Chips Chapman  
velomac@mindspring.com

### TMUG Committees and their Chairmen

*Programs:* Allen Emory

*Membership:* Alec Whittaker

*Web site:* Melanie Crain, pro tem

<http://www.tmug.org/>

Apple, the Apple logo, the MacOS logo and Macintosh are trademarks of Apple Computers, Inc., registered in the U.S. and other countries.

## DECEMBER HIGHLIGHTS

### Welcome to tmug!

David Parry, Anthony Tinkle, Tom Edelson & Donald Hughes



## Reality and Digital Pictures

by Charles Maurer

People often ask me if I think digital photography is as good as film or will ever become as good as film. I reply that for all but a few special purposes, digital is better already. Technically, my digital photographs are at least as good as the best conventional photographs I ever took with 2-1/4" x 3-1/4" (6 cm x 9 cm) film, and pictorially they are better. With my digital camera I can take pictures in the street that used to require a studio.

In this article I shall explain what digital technology can do that conventional photography cannot — how computers can produce more naturalistic pictures, not how they can produce special effects. To do this I'm going to start with perception, pass through art, and enter computers by the back door. Although this is an unusual route, it approximates the way I think when taking a photograph and it provides the only way I know for negotiating the maze of manipulations offered by photo editors. Although I shall mention some specific products (all of them available for the Mac as well as Windows), I shall not describe any in depth. The difficult part of digital photography is figuring out what must be done in the computer and which application can do it. Knowing that, it is rarely difficult to figure out how to make the application do its job.

This article is illustrated with a number of pictures. To see them appropriately, your monitor ought to be in rough calibration. If you have never calibrated your monitor, I suggest that you do it now. It takes about two minutes. Open the Displays preference pane, click the Color tab, click Calibrate to launch the Display Calibrator Assistant, select the Expert Mode checkbox, and then follow the instructions. When you come to the screen asking you to set the gamma, select 2.2.

*(Editor's note: I have added the pictures to this article, but I would still recommend monitor calibration. )*

For one reason that will become clear, I find some version of Photoshop to be necessary. For this reason I shall assume its use as a photo editor, although you need not own it to understand the article. Along the way I shall mention the differences among the last three versions (CS, CS2 and Elements) that matter for my approach.

### Eye vs. Camera

To begin with, let's dispel the notion that a camera records what the eye can see. It does not and it cannot because a camera functions nothing like the eye. With a lens of normal focal length, a camera records an image with a diameter of approximately 45 degrees. It records the entire image at once and the image ends up as a print with a range of intensity from black

to white of approximately one hundred to one. In contrast, the eye sees an area about 180 degrees across but it sees most of this with acuity that ranges from bad to dreadful. It sees sharply just in the central 1 to 3 degrees. To see a scene clearly, the eye must scan it and the brain must assemble the accumulated information. However, the eye rarely has time to sample more than small portions of a scene with its spot of clear vision so most of what you see has no optical source, it is an inference. Your brain infers information largely by generalizing from what it has encountered before. In doing this the eye and brain have to handle contrasts of light that exceed one million to one.

In short, when you look at a snapshot you took at the beach, the limitations of the camera mean that three-quarters of the scene will have been lopped off, the range of tones will be compressed tenthousandfold, and the information that remains will never be what you saw. Any appearance of realism will be an inference informed by learning and shaped by convention. It is not realism but verisimilitude.

Photographs may seem realistic but the technology of film prevents escaping photographic conventions, which are actually quite limiting. Less limiting is a paintbrush. A brush can produce every effect a camera can plus a great many more. Before photography, skilful and observant artists spent millennia working out how to represent reality on flat surfaces using this superior tool. Their work forms the most complete guide available on realistic ways to put pictures onto paper.

Most artistic techniques cope with two basic problems, problems that reflect the architecture of the visual tissue of the brain: how to imply something about form and space using (1) areas of brightness and (2) lines. These problems are not discrete and isolated any more than the tissue of the brain is, they are two sides of the same coin, but it will simplify our thinking to make a fuzzy distinction between them.

### Contrast

The eye does not see light per se, it sees changes in light — contrast. If two objects do not contrast with one another, to the eye they meld into one. This fact makes controlling the contrast of adjacent details to be paramount in importance. However, the real contrast of any scene can rarely be reproduced. As I said, the range of reflectance from the lightest to the darkest objects in a scene is rarely less than one thousand to one and often exceeds one million to one, yet the range of reflectance of pigment

*continued on page 4*

*continued from page 3*

against paper or canvas is approximately one hundred to one. On the other hand, even within a high-contrast scene, small areas can have very little contrast indeed.

From contrasting tones the brain infers three-dimensional objects. It does this through association, by matching patterns it has encountered before: a bright spot is a source of light, brilliant yellow may be fire and hot, areas that are darker tend to be removed from you or from light, bright areas tend to be near you or near light, tiny highlights on a face indicate sweat and heat, etc. To paint realistically, painters use associations like these to create optical illusions. This is easy because the eye scrutinizes only tiny areas at a time, so the brain cannot easily compare colours and tones across broad distances. As long as adjacent tones vary naturally, distant tones can be impossible optically yet still look right. You can see this in Rembrandt's painting of *Belshazzar's Feast*, linked below. The main source of light on the faces appears to be the writing on the wall, yet it is no brighter than the faces. It is not white but fiery gold, yet it is so far away from his face that nobody notices the optical absurdity. Also, with writing on the wall as the main light, the secondary light reflected off the invisible wall on the left ought logically to be much dimmer than it is.

<http://www.tidbits.com/resources/809/BelshazzarsFeast.jpg>



In other parts of the painting Rembrandt increased contrast where he had to maneuver within too limited a range to limit himself to variations in brightness. Look at the woman's red dress to see an example. Not only do the folds look three-dimensional overall, each tiny portion of every fold looks three-dimensional, even if you restrict your eye to small areas, areas where there is little difference in brightness from highlight to shadow. Every tiny part of the dress contrasts with the part

adjacent to it. Rembrandt could do this because he did not vary brightness alone, he independently varied hue and saturation as well. If you open the picture in Photoshop and set the Info window to HSB, you can move the mouse around and see some of this variation that has survived the miniaturization of the painting. (The real thing, which somebody long ago trimmed to a smaller size and different angle, is 66" by 82" or 167 cm by 209 cm.)

Filmmakers and commercial photographers create realistic photos similarly, by "cheating" lamps that are put on the set as props, lighting the set so that the light seems to be coming from those props. An example is the picture of the blacksmith at the link below. A logical analysis shows that no illumination can have come from the fire, but the eye is not a logical analyser. However, cheating like this takes more time than cheating on your taxes, especially in a still photograph where the illusion does not flit past your eye. That photograph took me a day to plan and a day to execute. (Among other things, I needed to wrap the entire workshop in aluminum foil, to prevent light from coming through chinks in the walls.)

<http://www.tidbits.com/resources/809/Blacksmith.jpg>



On the other hand, equivalent results can often be obtained without cheating by using a good digital camera and re-balancing the light digitally. An example is the dyer in the picture linked below. The version on the right shows the scene as film would have caught it; the version on the left shows it as it felt and as I remember it to be. It is probable that before I took the

*continued on page 5*

*continued from page 4*

picture, I noticed that the room light was bluer than the firelight — I do tend to notice such things — but my overwhelming perception was overwhelming heat and that heat is what I wanted to portray. To the visual system, so many cues to heat are present that the firelight in his face looks natural although it's logically absurd.

<http://www.tidbits.com/resources/809/Dyer.jpg>



*continued on page 6*

continued from page 5

The next example shows a more ordinary picture. The image on top shows what the scene looked like: a brightly lit bush in the foreground with a jungle of trees in the hills behind, gradually diminishing in size and clarity. However, although my brain perceived the bush to be bright, it was actually dark compared to the sky and the jungle was even darker. The scene presented a range of tones that nothing man-made can come close to reproducing. My camera's sensor "mechanically" compressed those tones into the image on the bottom. Slide film would have done the same. To make the picture look more realistic, I brightened the bush in the foreground and painted contrast into the jungle by varying saturation and brightness independently from each other and from hue.

<http://www.tidbits.com/resources/809/Jungle.jpg>



To manipulate contrast in this way requires three things:

- \* Capturing the information that you want to bring out;
- \* Making that information visible by lightening shadows and/or darkening highlights; and
- \* Adjusting colour not to make it look accurate — that is impossible — but to bring out whatever contrasts are necessary to make it look right.

To meet the first requirement, you need a raw, unprocessed image (not a JPEG) from a camera that can record a broad range of contrasts. In today's market this means a single-lens reflex camera. (For more information, see the "Image Quality" section of my article "*Picking a Point-and-Shoot Camera: Panasonic DMC-FX7*" in TidBITS-783\_.) When I convert the file to a standard format (I prefer the generic TIFF to Adobe's PSD), I set its levels of tonality to run the full extreme from black to white, with the middle set to look as good as possible.

<http://db.tidbits.com/getbits.acgi?tbart=08136>

Lightening shadows and darkening highlights comes next, with Adobe's Shadows/Highlights control. Photoshop defines shadows and highlights as dark or light areas larger than a certain number of pixels across. CS, CS2 and Elements all enable adjusting the amount of lightening or darkening but CS and CS2 also enable adjusting the size of what Photoshop sees as a shadow or highlight. I find that adjustment to be very important, and I use it for maybe one photo in three.

(Most of what Adobe left out of Photoshop Elements I do not care about — Elements is already more complex than it needs to be — but I found this one adjustment almost reason enough by itself to forgo Elements for the full Photoshop. The other reason is that Elements has limited facilities to handle 16-bit colour. Although 8-bit colour is usually sufficient, pulling apart tonality often requires finer intermediate colours to be present.)

Now look at the Rembrandt picture again, at the detail of Belshazaar's cape (page 4). The detail stands out because it is formed by brush-strokes with extremely high contrast from one to the next, extremely high local contrast. I make detail stand out in a photograph the same way by using an incidental feature of PictureCode's Noise Ninja, which is primarily a noise-reduction package (and one of the best). This feature is a slider that enhances local contrast. I often use it by itself without any noise reduction at all.

<http://www.tidbits.com/resources/809/BelshazzarsFeast.jpg>  
<http://www.picturecode.com/>

Now comes the paint. If an artist wants to adjust a colour on his canvas, he may change its hue, or he may daub on spots of complementary colours to reduce its saturation, or he may add some black or white touches to reduce or increase its brightness. With digital photographs I want to do the same. The product that enables me to do this is Asiva Shift+Gain.

<http://www.asiva.com/>

continued on page 7

*continued from page 6*

Shift+Gain is a Photoshop plug-in that lets you select areas or lines (useful to remove colour fringing) by any combination of hue, saturation, and brightness, and then alter those parameters individually. No other product can do this, except for a stand-alone package from Asiva that is too slow to use. Indeed, incredible as it may sound, Asiva has a U.S. patent on this approach to manipulating pictures.

Shift+Gain works differently from any other application and took some time to understand. However, although it was confusing at first, it soon came to seem simple. To accomplish in Photoshop most of what I do in Shift+Gain would require far more skill and patience than I can supply.

I find Shift+Gain to be an indispensable tool for digital photography - the only indispensable tool, the only tool for which I do not know of any functional equivalent. Unfortunately, it will not work in any application other than Photoshop, not even applications like GraphicConverter that can run most other Photoshop plug-ins. It is compatible with any recent version of Photoshop, but it does require Photoshop, which is why I am ignoring possible alternatives to Photoshop in this article.

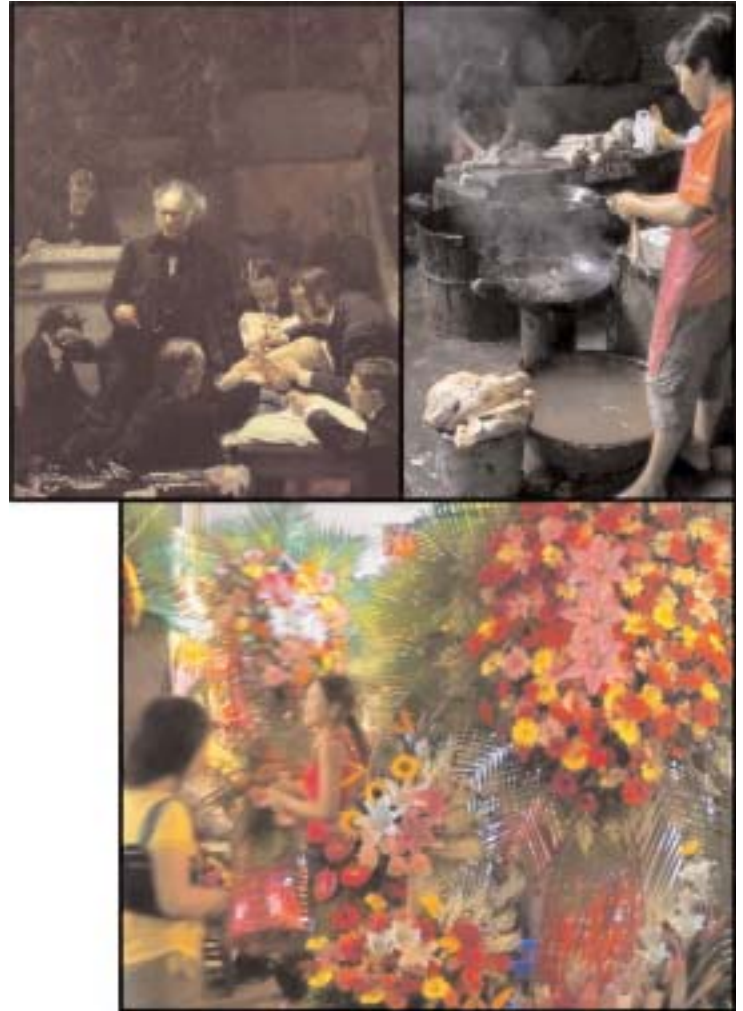
Those three sets of tools can handle nearly all the manipulations of contrast and colour that I have had any need for: (1) the controls in Photoshop CS/CS2 for levels, shadows and highlights, (2) the local-contrast control in Noise Ninja, and (3) Asiva Shift+Gain. Occasionally I also use one of Asiva's other plug-ins, which work similarly but do slightly different things. I have found that Asiva's plug-ins, combined with Photoshop's basic selection tools, obviate the need for masking to achieve ordinary pictorial effects.

Only one of Photoshop's colour adjustments do I find to be particularly useful. Sometimes, after I have adjusted the colours to bring out contrasts, the picture shows an overall tint. Now, no tint exists on its own, a tint is merely an offset from a standard of comparison. In a photograph, the eye's standard is usually a pure white highlight or the paper's margin. If a neutral white or grey looks coloured in comparison, then we see a tint. Removing a tint is usually a simple matter of shading the picture just enough to neutralize that white or grey. Every other colour changes a bit, but the contrasts among them will remain. It's difficult to remove a tint manually because the brain adapts so readily to changes in colour that a wide range of adjustments seems okay until you print out the picture. Photoshop can remove a tint mechanically; the mechanism is hidden in the Match Color command.

One final consideration about colour comes with dim light. In sunlight we see in colour; in moonlight we see in monochrome; in transitional "mesopic" levels of dim light we see

partially in monochrome and partially in colour. When painters want to represent dim light, they portray it mesopically. You can see this with the musician at the back of the Rembrandt and you can see it even better in the Gross Clinic by Thomas Eakins, the picture on the left at the link below. The students in the shadows are nearly monochromatic but the monochrome contains hints of colour, often quasi-random streaks and blotches. (Note that the original painting is 96" by 78" or 243 cm by 198 cm.)

<http://www.tidbits.com/resources/809/GrossAbattoirFlowers.jpg>



Film does not portray dim light in this way, nor do most digital sensors, but the Foveon sensor does. (See "Sense & Sensors in Digital Photography" in TidBITS-751\_ and my followup for a discussion of sensor types.) Film and digital sensors generate low levels of granular noise. When a normal amount of light strikes the film or sensor, the noise is usually hidden within the image, but when little light strikes it, the noise becomes more evident. At some dim exposure to light the image disappears within the noise: that defines the limit of sensitivity.

*continued on page 8*

*continued from page 7*

The random dots of this noise can be smoothed over but detail becomes smoothed over with them and at the limit of sensitivity, all detail disappears. However the Foveon image sensor works differently so its granularity looks different. The Foveon shows fewer specks but replaces them with intrusions of incorrect colour. At first this reduces saturation then, at the lowest levels of sensitivity, it causes random streaks and blotches.

<http://db.tidbits.com/getbits.acgi?tbart=07860>

<http://db.tidbits.com/getbits.acgi?tbart=07906>

Reduced saturation and random streaks and blotches of colour are exactly the techniques that artists use to represent dim light, and the Foveon's noise can be used to do the same. I smooth out the granular noise with Noise Ninja — there is rarely so much of this that Noise Ninja loses any detail — then I use Shift+Gain on selected areas to control the discolouration. My goal is sufficient discolouration to add contrast for the eye but not so much as to be noticed. You can see the effect in the Chinese abattoir to the right of the Gross Clinic painting you just loaded.

Do note, though, that desaturation and blotchiness are not the norm in Foveon photos. They are normally hidden in depths of black and become evident only if you bring them out by pushing the sensor to its limits. More normal is the picture of the flower market — the third one on the page. I took both pictures indoors and exposed them at ISO 1600.

## Perspective

So far we have been talking about how to represent space using tonality, now let's shift to representing space using lines. This is the problem of perspective.

During the Italian Renaissance, artists worked out a geometry of linear perspective, geometry that appears superficially to fit perceptual norms. In fact, however, it does not. The "laws" of linear perspective need usually to be broken, else the picture will look wrong.

The laws of perspective dictate that parallel receding lines converge. They converge if they are receding horizontally like railway tracks and they converge if they are receding vertically like skyscrapers seen from the street. But consider vertical perspective. If the angle of view portrayed is only a little bit upward, then your brain may not infer that objects are converging at a distance above you, your brain may infer that the objects are not plumb. Of course, if those objects are walls of buildings, then your brain concludes that they are not falling inwards, for just as you assume that boards are straight, so do you assume that walls are plumb. However, for the same reason — because

you assume that walls are plumb — buildings look more natural when all the vertical lines are upright and parallel. You can see an example of this issue in the two images of the temple pictured at the link below. A correction like the top image with film would have required the careful adjustment of a view camera on a tripod but it took me two minutes in Photoshop. (Elements or CS can fix perspective but CS2 makes it easier through a new Lens Correction item in the Filter > Distort sub-menu.)

<http://www.tidbits.com/resources/809/VerticalLines.jpg>



The same adjustment is useful for horizontal lines. When horizontal lines converge, buildings can appear to be constructed on a hill and roofs can seem to have unusual inclines. To minimize ambiguity, vertical lines ought to be plumb and horizontal lines ought to be level unless the reason for them not to be is obvious. Clear verticals and horizontals provide a frame of reference that lets oblique lines stand out.

Pictures of buildings obviously benefit from this approach, but often pictures of people do too, although more subtly. You can

*continued on page 9*

*continued from page 8*

see an example in these two pictures of children, linked below. The picture on the top is stronger because the children are sitting on a level platform, not a tilted one.

<http://www.tidbits.com/resources/809/Children.jpg>



In fact, the laws of linear perspective need to be violated even when photographing something straight on. If you look straight at a picket fence or a wall of bookshelves, an optically correct perspective would have the lines of the fence or bookshelves converging both to the left and to the right. This would look so silly that nobody would paint them this way. For the same reason, camera lenses are corrected to distort linear perspective so that a rectilinear object casts a rectilinear image.

This presents an interesting problem that can be solved with a brush or computer but not with film. The farther out from the

centre an object extends, the farther its lines will be pulled apart and thus the more it will be enlarged, yet objects in the centre will never be enlarged, distorting relative sizes. The wider the lens's angle of view, the greater the distortion. This distortion can be seen with any wide-angle lens and becomes disproportionately more severe the wider the angle of view. When straight lines are not involved — in many landscapes — it often looks more natural when relative sizes are maintained at the expense of convergence. This can be approximated in Photoshop CS2 by adding convex “barrel” distortion, a distortion that reduces the rectilinear correction of the lens.

(Note that only CS2 offers that control. CS2 also makes it significantly easier than CS or Elements to correct converging and tilting lines, once you find the new controls. In CS2, all of the lens corrections are buried under Filters > Distort, although File > Render still shows the subset of corrections that is shared with CS and Elements.)

Of course, adding convex distortion is unacceptable if straight lines are involved. A certain amount of convex distortion may not be noticed in landscapes, but curvature stands out absurdly in pictures containing buildings. An alternative fudge is to squeeze the picture from the sides. To do this I use a \$20 Photoshop plug-in called Squeeze.

<http://www.theimagingfactory.com/>

I also ought to mention the portrayal of depth through having only one plane of the picture in focus. This effect can be achieved with a brush, but it rarely is, because it does not mirror what the eye sees or the brain perceives. The eye sees only tiny spots sharply, and it sees tiny spots wherever it looks: from these the brain perceives infinite depth of field. To control attention and suggest different qualities, a painter will vary the softness of edges across a picture, but this variation is much more subtle than a mis-focussed lens.

To vary hardness and softness within a picture, I used to use a view camera that allowed me to tilt and swivel the lens, and I varied the character of the light. A digital camera makes this a lot easier. My digital camera usually provides infinite depth of field with no special measures and I can use digital techniques to control softness like a painter, as I did in the flower market example previously shown. The flowers just behind the smiling

*continued on page 10*

---

*continued from page 9*

girl are soft, but the ferns behind them are sharp, as is every other object in the picture except for the woman moving into it.

This was possible for two reasons, both tied to the camera's image sensor. First, the ISO speed of negative film is based on the least exposure necessary for acceptable snapshots. To extract high quality usually requires doubling the metered exposure. In contrast, to extract the best quality from my digital SLR, I usually halve the exposure. That is two f-stops' difference, which represents a lot of depth of field. On top of that, the sensor in my camera is smaller than 35mm film, which means the same f-stop gives more depth of field. The difference is 1-2/3 stops. Thus, for any given amount of light, I obtain nearly four f-stops' more depth of field than I would get were I shooting 35mm negative film.

When everything is sharp within a photograph, photographic compositions open up. People don't just look at my pictures, they look inside them, combing them for detail — and they find it because I have controlled the details' contrast. With so much information to look at, my 8" x 10" (A4) printer seemed too small.

Finally, to finish up my comparison of the various versions of Photoshop, I ought to mention two new features of CS2 that are useful for preparing enlargements, a "spot healing brush" and "smart sharpening." The former I find to be a modest but significant convenience, but the latter is an important feature. It tightens up a lens's inescapable spreading of points into blurry circles, and it reduces blur from movement. In my mind, this feature combined with CS2's improved distortion controls makes the upgrade from CS worth the purchase. I detest a Windows-like copy-protection scheme that Adobe have begun to employ — it prohibits the fair use of your purchase if you work in different locations — but I swear at CS2 less often than I did at its predecessors because it permits me to hide from sight the vast number of menus that I never use and to edit or remove keyboard shortcuts. With CS2, no longer do windows fly about the screen and change their colour because one of my fingers inadvertently touched a key.

*PayBITS: If you found Charles's discussion of visual perception and digital pictures useful, please support Doctors Without Borders: <http://www.msf.org/msfinternational/donations/> Read more about PayBITS: <http://www.tidbits.com/paybits/>*



## RAID Backups

by Tom Boucher, *tmug*

**Question:** Does anyone use RAID? I don't quite understand the purpose or how to use it, though I read some about it.

**Answer:** With more and more of my life going digital (Photos, Music, Movies, and more) I wanted some way to protect my stuff from simple hardware failures. Meaning backups are great but when you have to restore hundreds of GB of data DVDs get old rather fast. RAID has been around now for at least 15 years that I can remember. My day job includes working with large Windows/Linux/Unix and what-not servers.

The idea of RAID is an algorithm to write data across multiple hard drives to get redundancy in the event of a drive failure. There are three common levels of RAID.

**RAID0 = the 0 in RAID0 equals zero protection.** RAID0 takes two (or more) drives and makes them look like one big hard drive. If any drive breaks, data is toast. However with more drives working on your reads and writes the guys that love high performance love this because they get better throughput to their disks.

**RAID1 or Duplexing.** This is taking one drive, and making an exact copy of it. Every time you write to one, the other is written too exactly the same. While one of the more resilient methods it's also the most expensive in for every single drive you want to protect you have another drive backing it up. However this is also the highest performer with protection.

**RAID5, or Distributed Data Guarding.** This method requires three or more hard drives. Every time you write a piece of data to this array two pieces are written to two drives, while a third is written a math problem that can solve for the data on any one of the other drives. So you can lose a hard drive in this group and still keep on trucking. Lose a second drive before

the one is replaced though and data is gone.

For some time I had a network server that had four 250GB Drives 'RAID5' protected that ran linux. This was pre-Xserve RAID days. However the controller I was using (Adaptec 2400A) failed to mention when a drive had failed and then a second one failed after a power failure and I lost all my photos from Easter 2004 and a family vacation to Wilmington we took because I hadn't done a recent backup within the last month.

Since then I've gotten rid of that solution and I use external drive for the daily backup using Backup 3.0 from .Mac and then weekly burn to DVD-R using my Superdrive.



*Example of Apple's Xserve Raid. ©Apple.*

It works because I just lost the internal drive in my year old iMac G5 and spent two days restoring about 60 some DVD-Rs. While not the speediest of restores my stuff was protected and I lost nothing this time (whew).

Hope that helps.

—Tom