
Does HDR Mean *You No Longer Need Lights* When Shooting Digitally?

by **John Siskin**
and **Richard Lynch**



Figure 1. The finished strobe version of the image, which used five strobes.

Architectural shots are an important part of my business [John Siskin reports the first part of this article]. When I do a shot with strobes, I often work for a couple of hours and move a couple of hundred pounds of equipment in the process. That is a pretty normal way of doing architectural lighting for me. The newer HDR (High Dynamic Range) feature in Photoshop, introduced with Photoshop CS2 and improved with CS3, has potential for making my job easier by reducing the need for all that equipment and automating the process of merging multiple exposures. The purpose of HDR is to make a high-bit-depth composite that does a better job of retaining highlight and shadow details than a single shot can. I decided to see what could be done using natural light and HDR and how that compared with a shot in which I set up strobes.




Figure 2. This shows the image with existing light only. Exposure is on automatic.

I talked about my architectural photography with Richard Lynch, a Photoshop expert who teaches with me at BetterPhoto.com. We decided to take an HDR shot of an architectural subject without lights and compare it with a version of the same shot done with strobes (Figure 1).

Setting up the test

The first thing to do when shooting an architectural subject is to establish the camera's location. I used a Kodak DCS Pro 14n, which I like because I can use my Nikon lenses and still get full-frame capture. One problem with this camera, however, is that it doesn't perform well at an ISO greater than 160. This means I use a considerable amount of light when I shoot architectural subjects, though it's a lot less light than when I worked with a 4x5 view camera. In this case, I actually was looking for a place with difficult or impossible lighting to better do the test.

I chose the camera's spot based on the windows and the sunlight falling to the left of the camera. I actually waited a while for the sun to bring that light more into the frame. My other consideration was the placement of the front door. I liked this location, but I could have set the camera a bit higher to separate the door from the banister. You can see what the shot looked like on auto exposure in Figure 2.

I decided to first capture the image using strobes. Strobes give you control over the contrast in a shot in a couple of ways. First, you can fill in your shadows by adding more light. Second, you have control over the highlights because the shutter affects the ambient light without changing the light from the strobes. The strobes have such a short duration that the shutter doesn't affect light from the strobes (see my article in *PHOTO Techniques*, July/August 2007 issue for more on this).

I used five strobes, plus some daylight, to light the strobe version of the shot (Figure 1). This was a case where I didn't need a lot of power, but I did need light in a lot of places. Just as you wouldn't normally use one light to illuminate an entire living room, you wouldn't use just one strobe to light a room.

The goal in a shot like this is to bring the contrast down. The unlit version of the room (Figure 2) has too much black, too



Figure 3. This shows the position of the largest strobe, the computer, and the camera. The strobe and the camera are close together to provide a similar angle for the light and the lens.




Figure 4. This diagram shows how the lights were set up at the location. Strobe D had about 300 watt-seconds; strobes A and B were set at 100 watt-seconds; strobes C and E were set at 200 watt-seconds.

much white, and not enough midtones. If you were going to fix this problem for just sitting in the room, you might use lamps with lampshades. The shades effectively make the light source larger and softer. To do the same thing with a strobe, I usually use an umbrella. In this shot I used a 60-inch umbrella with a Calumet 750 Travelite.

I like to start with the biggest, broadest light, so I placed the Travelite and umbrella on the second floor near the camera, where the light covers most of what the camera sees (they have a similar point of view). I used the Travelite, which is a pretty powerful light, because it has continuously variable power output. This makes it easier to bring the level of strobe and daylight into the right relationship. I used this strobe at about 300 watt-seconds.

You can see the setup in Figure 3. The camera is tethered to my laptop, which is a big help. First, I can see the image on the laptop—I love this ability to see the image, at a decent size, a few seconds after I fire the shutter; it's critical for architecture. I used to go through a huge amount of Polaroid doing shots like this, yet the image on the laptop is better. Another nice thing is that I can fire the camera and change the aperture and shutter speed from the computer. This reduces the chance of



Figure 5. This shot has four of the five lights finally used. You can see that the left side of the round room and the ceiling are too dark.



Figure 6. This is the snoot made from a piece of plastic tubing. The light it gave from the side and the top worked out very well.

moving the camera—which is particularly valuable when I shoot the several exposures I’ll use for the HDR file. Even so, I have a bag with a gallon bottle of water hanging from the tripod. This helps to keep things stable. I like to use a water bottle for extra weight since I don’t have to fill it up until I get to the site. That makes carrying the equipment a little easier.

I next added three more strobes; the placement on these was pretty obvious to me. These lights are filling in parts of the shot that would otherwise be in shadow. I needed to have light in the bottom of the cylindrical entry, opposite the front door, so I added a Norman 200B strobe. I spread this light with a 45-inch ribless umbrella. I like ribless umbrellas because, with the bounce fabric in front of the ribs, they spread light more evenly and their reflections are cleaner. I also used the same setup in the ground-floor hall to the left of the camera. Both these lights were set at 100 watt-seconds. I used another Norman 200B in the bedroom on the right side of the camera on the second floor. With this light, I used a 30-inch white shoot-through umbrella, which spreads light behind and in front of the strobe. Consequently I don’t have to be as critical with light placement. Remember that I only need to be able to see into this room, not be able to see details. (You can see the placement of these lights in Figure 4.)

I like Norman 200B strobes for architectural work for a couple of reasons: first, I don’t need to find wall sockets for power since these units are battery powered. The batteries do add some weight to my kit, but they save me from stringing power cords everywhere. These strobes also pack a lot of power into a small package—200 watt-seconds, enough for a lot of lighting applications with a digital camera. Finally, I have a lot of accessories that fit the Normans.

At this point the image looked like Figure 5. Not a bad start, but in the left side of the tower the stairs descend into shadows, and the top is not light enough. This is a significant problem, and brings up one of the biggest dilemmas with

architectural lighting: you can’t always put a light where you want, because it may be in the shot. When you look at Figure 5, you can see all the places I couldn’t put a light. The stairs were my best choice. So I put one more 200B on a very small light stand made from a vice grip. A piece of white plastic plumbing pipe on the 200B served as a homemade snoot (Figure 6), which put a lot of light out the front—in this case toward the ceiling—and some light out the side. I have gotten lucky with this tool on several occasions, and I don’t mind admitting that it performed better here than I thought it would. I was pleased with what happened to the side and the top of the tower walls (Figure 6).

I spent a little more than two hours working on the shoot, and I still needed to clean up. So, as planned, I took a series of shots for the HDR file in the same place, using only ambient light. I made exposures from five stops under normal to six stops over normal. So that everything was in register, I did this with my camera tethered to the laptop and made all my exposure changes without even touching the camera.

I’ll let Richard take over from here.

HDR without lights

John forwarded me 13 exposures for the HDR shot. I wanted to compare an HDR version of the image to one of the images converted in Adobe Camera Raw, and a shot that opened and combined several of the exposures manually. I dove right in and opened the shots in Photoshop CS3 as an HDR image. To do this you choose Merge to HDR from the File > Automate menu. A dialogue appears (see Figure 7) that allows you to choose the images to include in the merge. I chose all 13 so I’d be able to make the most of what John shot.

As John had taken precautions with his equipment, alignment was virtually unnecessary. The next step was to look at the image in the Preview window (see Figure 8), which shows all the exposures in a listing at the left along with a preview of

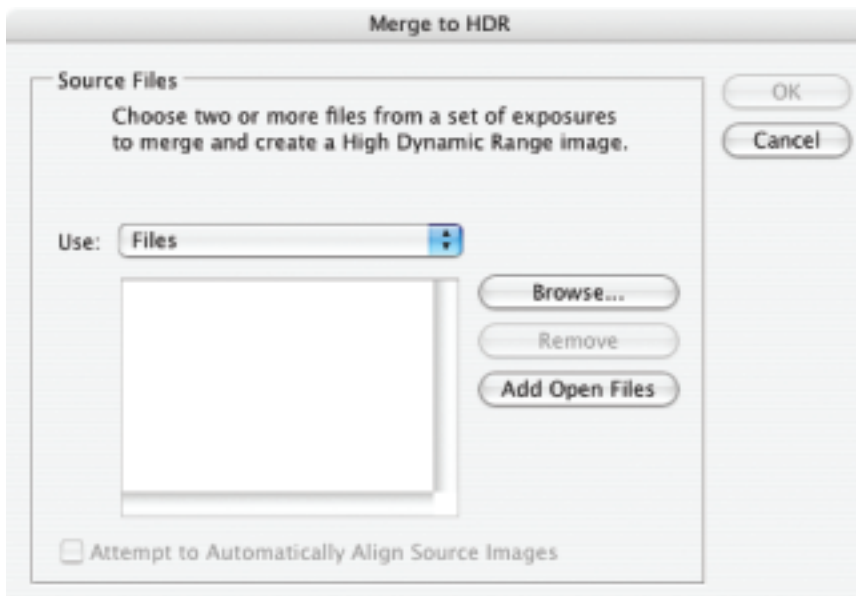


Figure 7. The Merge to HDR dialogue box.

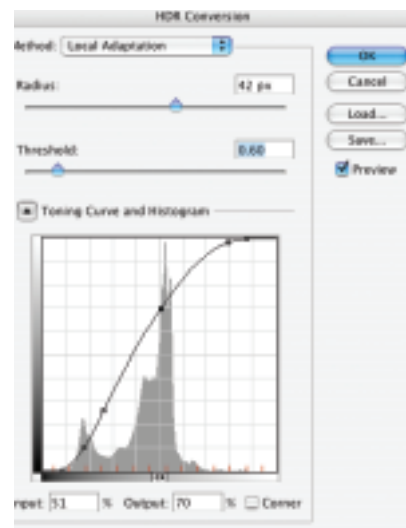


Figure 9. The dialogue for reducing the file from 32- to 16-bit depth.

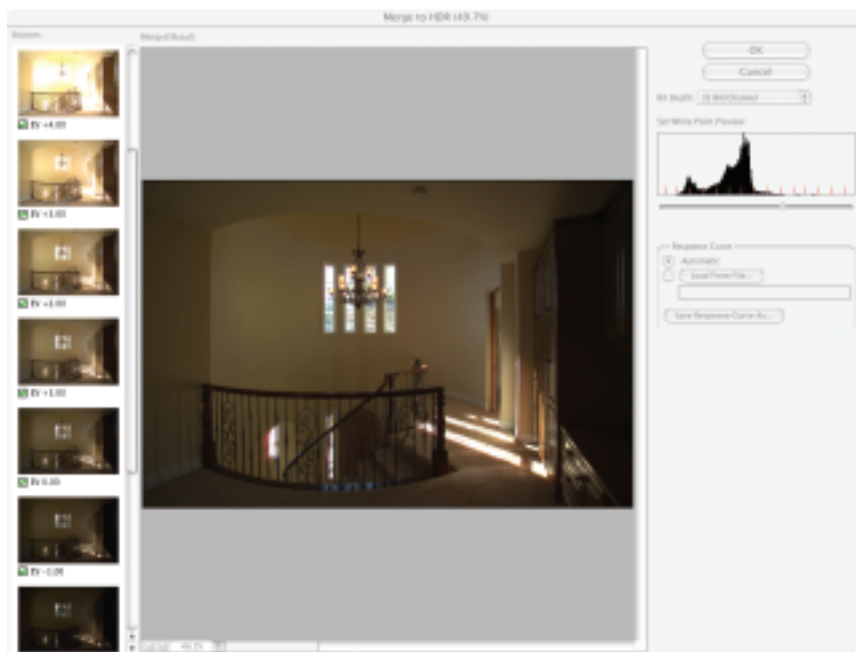


Figure 8. The image in Preview; note that all the single images are visible on the left of the screen.

the result. To the right of the screen, you have a single slider for adjustment that essentially lets you brighten or darken the overall response to the merge of the exposures. You can remove exposures from the composite at this stage by unchecking the box below the thumbnail views, to the left of the screen.

The merged version of the image does not look all that different from the best exposure of the bunch converted carefully from Raw. The key advantage is the 32-bit information, which leaves you with the potential of about 4.7×10^{21} times the color of an 8-bit image. The reality, however, is that you can't use all that image information either on-screen or for output; these, so far, are 8-bit arenas. The extra information also balloons the file size.

While you can make adjustments to the image at this stage, the options are limited in Photoshop, as all the functions available for 8- and 16-bit images are not enabled for 32-bit. The next step is, logically, to reduce the 32-bit monster into an easier-to-manipulate 8- or 16-bit image. When you convert to another mode (e.g., Image > Mode > 16 Bits/Channel), Photoshop brings up a conversion dialogue that allows you to make some adjustments for the conversion (Figure 9).

While the dialog can be helpful in getting closer to what you might expect, and the Curve gives you the appearance of having greater control, the adjustment you are making at this point is essentially a linear one in that the color and tone are affected across the entire landscape of the 32-bit image. A more satisfying conversion would allow you to enhance the influence of one or more of the exposures within a range; in other words, bring up the highlight information from the highlight exposures only—or even from one specific exposure.

The natural result of conversion to a lower bit count is often called “decimation,” and here the result is not all too different from the inference of the word: you give up enormous potential information. You also have to endure the reality that enhancing detail in the highlights and shadows necessarily compresses the global contrast in the image. In reality, the science of it can lead to only a marginally “better” image than a best single exposure. Please see a comparison of best exposure Figure 10 and the HDR version in Figure 11.




Figure 10. This image was made from the best single exposure.




Figure 11. This is the HDR version of the shot.

The problem here stems from the desire to make a realistic rendition of the scene coincide with what is essentially an artistic vision of it. Lighting will help to accent specific areas to paint the scene and create distinct and controlled emphasis, whereas multiple exposures of the same scene with the same light can, at best, reproduce that same scene with somewhat more detail in the highlight or shadow—compromising the midtones and overall contrast. It can produce some interesting effects if used more like a filter than a tool, but the reality is that the dynamic limitation of capture still remains.

A third version of the image can be created by manipulating the multiple exposures to isolate those areas that are most pleasing in Photoshop. Detail from the windows taken from one shot, the chandelier from another, detail from the dark portions of the cabinets from a third . . . all coming together in Photoshop with a great effort of masking, patching, blending, and layering. Figure 12 displays the result of some three-plus hours of work in Photoshop. The image still requires some work to match the result of John's effort in lighting. Here detail from beyond the windows becomes more defined, details in the room remain, glows and hot spots are reduced as desired, and selectively and mechanically the result is balanced and drawn together in a way that no tool in Photoshop can do on its own. The scene beyond the windows, for example, can be isolated from the rest of the scene and corrected for dynamic range and color in a way current HDR controls cannot imitate.

Conclusions

In all, we found that post-processing to create a good rendition of the image takes as much or more effort than working to light it right in the first place, but is a viable option when you either don't have the skills for lighting a scene or you lack the equipment. As with any successful result in Photoshop, the best outcome is produced not by applying an individual feature or function, but by working to achieve balance in the scene.

The tools in Photoshop, no matter how technologically advanced, are not magic; neither the computer nor the camera can see a scene and cannot make an artistic judgment as to what looks right. The only magic that can be had is from




Figure 12. This image is made manually by taking the best part of each file. A lot of work, but no cases of lights to carry around.

achieving your vision—which will likely come about more from using all the tools at your disposal rather than hoping one will quickly create the result for you. A photographer has to use both eyes and brain to decide how to use these tools. John says he doesn't think he has found the perfect tools yet; however, he did learn to use Photoshop to improve what he was doing with lights. He also thinks HDR works very well for large public spaces that might be almost impossible to light, and would be helpful in situations where he can't use strobes. ■

John Siskin is a commercial and fine-art photographer specializing in product images and portraiture, as well as macro and architectural photography. He has taught photography for more than 25 years. He currently teaches lighting and portraiture photography at BetterPhoto.com. His Web site is www.siskinphoto.com.

*Richard Lynch is the author of seven books on digital-image editing, including his most recent books, *The Adobe Photoshop Layers Book* (visit photoshopcs.com) and *The Hidden Power of Photoshop Elements* (visit hiddenelements.com). He does digital retouching for advertising, portrait, and headshot photography, and teaches courses in image editing with Photoshop and Photoshop Elements at betterphoto.com.*