

Using Photoshop CS 2's Merge to HDR to Simplify Complicated Exposures

by Josh Anon

joshanon@mac.com

How many times have you gone to capture a beautiful shot, be it a sunrise with flowers in the foreground or a scenic mountain shot with the sun only illuminating part of the mountain, only to discover that you can't capture the full dynamic range of the image with just one exposure. For years, photographers have used techniques such as split neutral-density filters to reduce the dynamic range of a scene so that they could take one image capturing what they saw. Although powerful, these tools are limited in very complicated compositions, such as a dark mountain with a bright sky, a moderately bright reflection, and a flower in the foreground. However, in Photoshop CS 2, photographers have a new tool, Merge to HDR, that can help capture what you really wanted to shoot!

Understanding Digital Images

Before we dive into how to use this tool, there are a few basic concepts about digital images to cover. Digital images are typically represented as combinations of three colors, red, green, and blue. A pixel is a combination of a red, green, and blue value (each of which is referred to as a channel) representing a point in the image. Generally, we say that each of these values range from 0 to 1. However, that implies that each channel has infinite precision because it can have fractional values (e.g. your images could distinguish between a value of .5 and .50000001), but digital files can't readily store that much precision: keep in mind that computers use a series of bits (0 or 1, two states representing on or off) to represent data. Digital images are typically either 8 or 16-bit. Eight-bit means that each red, green, and blue channel stores a round number between 0 and 255 (2^8 possible values) and 16-bit means that each channel stores a number between 0 and 65535 (2^{16} possible values). A pixel in an 8-bit image whose red, green, and blue values are (0, 0, 0) would be pure black, and one whose values are (255, 255, 255) would be pure white. Similarly, a pure red pixel would be (255, 0, 0).

The next concept is contrast in an image. Simply put, contrast is the ratio of the luminance of the brightest to the darkest points. In Photoshop, for both 8 and 16-bit images, the contrast is effectively the same (remember that they both just represent 0 to 1, black to white, but are just scaled); only the number of possible colors changes because a 16-bit image has greater possible precision.

32-bit Images

The obvious question is how can you increase contrast. The answer is rather simple, but it's conceptually a little hard to understand; allowing for values greater than 1 for white gives increased contrast. In Photoshop, 32-bit images allow for these greater values because when you create them from a series of images at different exposures, they initially have no fixed white point. Images that have a white point value greater than 1.0 and can have very large dynamic ranges, beyond what you can display on your monitor, are called high dynamic range (HDR) images.

To help understand, imagine that you take two pictures of the sun, one as a 16-bit image and the other as a 32-bit image. Let's also pretend you have two monitors--a 16-bit one and a 32-bit one (by comparison, most computer monitors are only 8-bits). White on the 16-bit monitor would be a brighter than your regular monitor, and both images would appear very bright but about the same brightness. Now if you turn to your 32-bit monitor, which is capable of displaying very bright whites, meaning there's no need to clamp your white values to 1.0, the 16-bit image would look roughly the same as it did on the 16-bit monitor, but the 32-bit image would be really, really bright. You've increased contrast by not clamping your white point to 1.0.

With digital photography and Photoshop CS 2, you can take a series of images of a subject at varying exposures, combine them into one HDR image, which you use as an intermediate step to your final image, containing all of the information from your image series. You then compress the 32-bit HDR image back into an 8 or 16-bit file that you can print and use like a normal image. This technique simplifies complicated exposures because you no longer have to worry about using filters to reduce the dynamic range of the scene without losing detail or making complicated masks to combine images. You also don't have to worry about what you're exposing for in the image, because you're exposing for everything!

Let's run through this step-by-step.

Step 1 is taking a series of images at varying exposures. Typically, you want to capture about three to seven images, ranging from roughly -2 to +2 at 1-stop intervals. If the subject is really dark or really light, you want to add another exposure or two on the appropriate side. Use a tripod! You'll want to have your images as similar as possible. Autoexposure bracketing can help you quickly take a series of images.

Step 2 is turning the images into an HDR image. Photoshop can process both RAW files and JPEG files.

1. From the File menu, under the Automate submenu, select "Merge to HDR...." A panel similar to Figure 1 will popup.
2. Pick the images, and check align images. Aligning images will correct for very small changes in your camera position between exposures.
3. Click OK. What's actually going on after you click "OK" is that Photoshop is determining the camera's response curve using the exposure information stored in the files. Put simply, if you double the exposure time for an image, the red, green, and blue values recorded will not necessarily double; Photoshop is determining how these pixel values change as the exposure changes. This need for different exposures to create the response curve also explains why converting a raw file twice--once at a brighter exposure and once at a darker exposure--and then trying to use it to create an HDR file fails--it's still just one exposure time underneath.

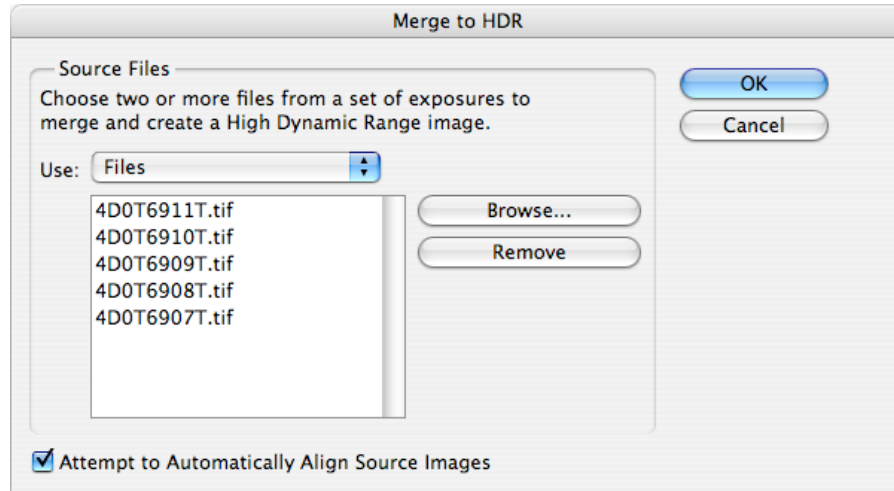


Figure 1: The Merge to HDR panel

Photoshop will return with a window asking you to set the white point, as is seen in Figure 2. This is where you set what you want the brightest point in the image to be. Everything brighter will lose detail and be clamped to white. Use the preview window and the slider to determine the brightest point in the image. Typically, it's the last point on the right where the histogram still has data. On the left, Photoshop shows you the relative exposure values of your different source images. You can uncheck some of your source images if you decide you don't want to use them to compute the HDR image. For instance, if you jiggled your camera slightly on the last image, deselecting the image will remove the ghosting that appears in the HDR preview, ghosting resulting from Photoshop not being able to properly align the image.

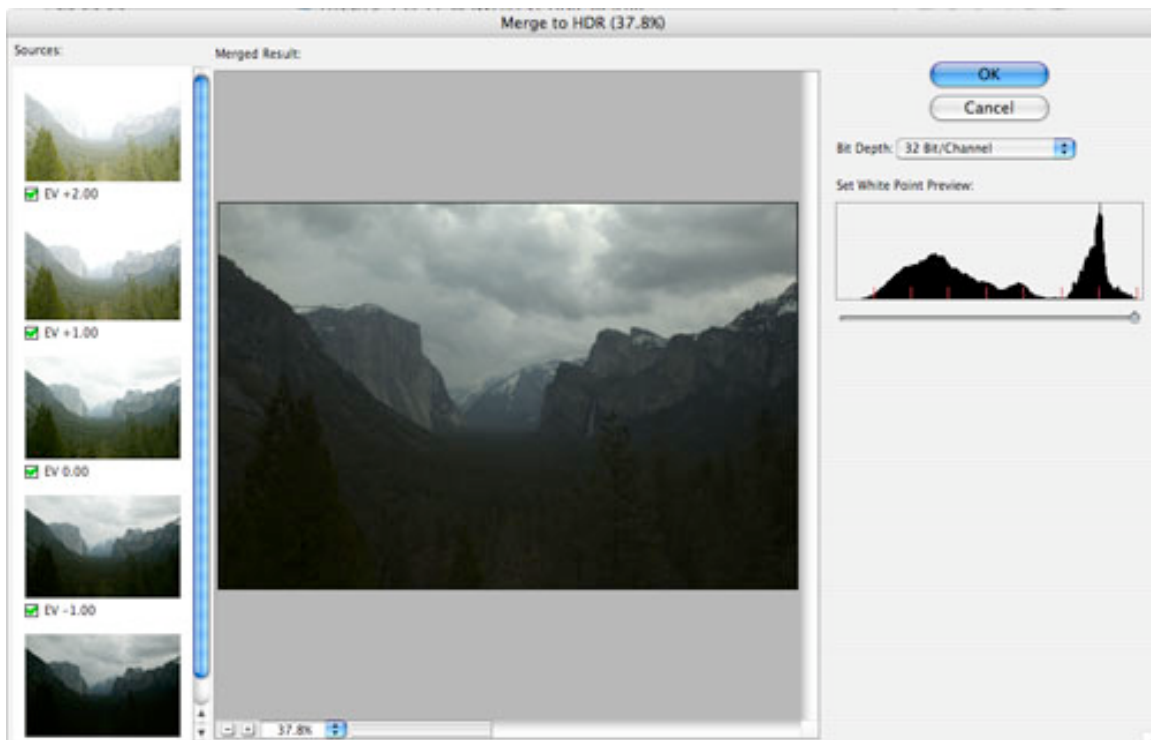


Figure 2: The white point preview window and the included exposures on the left.

At this point, if you were to click OK, you would have an intermediate 32-bit image. You could save this image, as it represents the entire image series you initially took, and try converting it different ways. The image might look odd on your screen--very dark, very bright, or very contrasty. It looks this way because your monitor is only capable of displaying a small part of the information in the HDR file. You can change what part of the information we're viewing by changing the tool to the left of the horizontal scroll bar to display a 32-bit preview control and adjusting its value. When you're ready to compress the image to an 8 or 16-bit image, select the desired bit depth from the Image/Mode menu.

Alternatively, before you click OK, you can change the popup from 32-bit image to either 8 or 16-bit. This simplifies things because rather than getting a 32-bit file and then changing it to 8-bit or 16-bit using the Image menu, you can go directly to your desired bit depth. Photoshop will automatically bring up the HDR conversion dialog after it has finished making the HDR file.

In the HDR conversion dialog (Figure 3), there are four different dynamic range conversion methods. Each applies the compression method to the pixels and then clamps the pixel values to the traditional 0 to 1 range.

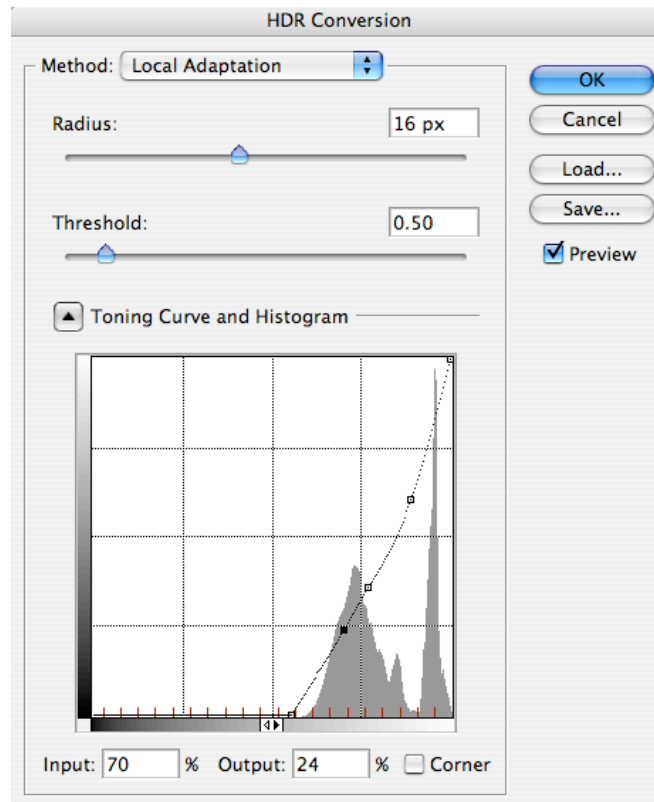


Figure 3: The HDR Conversion panel

The first method is "Exposure and Gamma." The exposure parameter basically scales the

pixel values, and the gamma value applies a gamma curve to the image. You can display the toning curve and histogram, but you can't modify it in "Exposure and Gamma." Some images look good by sliding the exposure value to the left (decreasing it) and sliding the gamma value to the right (also decreasing). This has the effect of first darkening the image so that more of the image is displayed and then applying a gamma curve to brighten up the midtones.

The next method is highlight compression. There are no options, and again you can't modify the curve. A highlight compression applies an operation that leaves the darker values roughly the same but scales the brighter values so that the brightest possible value is 1.0, compressing the highlights. This method tends not to look very artistic, although it might produce some interesting effects sometime and is at least worth looking at occasionally.

Equalize histogram is another method that does not give very artistically pleasing results, although it is useful for scientific images. Histogram equalization adjusts the pixels so that they cover the entire histogram instead of potentially having luminance values for which there are no pixels. There are no controls for equalize histogram, and it's generally not very useful for our purposes.

The last method, local adaptation is the method that gives you the most control. It's based on the idea that images can be divided into a base layer and a detail layer. The base layer has information about the large variations in the image (which Photoshop determines using a blurred version of the image), and the compression (defined through a curve) is applied to that layer. The detail layer contains the subtle tonal differences that make the image interesting, and it is combined with the compressed base layer to get the final image. The radius and threshold parameters in Photoshop control how the base layer is defined. More specifically, the radius determines how much the image should be blurred when determining the base layer, and the threshold parameter determines how large of a step between pixel values there needs to be to determine detail. You'll want to keep the radius moderately small (the default 16 pixels tends to be a safe place to start) so that the base layer looks roughly the same as the image. If you have lots of very fine variations in brightness, you can make the radius even smaller. Threshold is a bit tricky. If you make it too small, most of your pixels will be included in the base layer and your image will look flat. At the same time, too large, and your pixels will become part of the detail layer and your image will look noisy.

An approach that generally gives a great result is to use local adaptation with the default parameters but to modify the tone curve. Drag the left point (black) to roughly where there starts to be data in your histogram. Drag the right point (white) to roughly where there stops being data in the histogram. Then, adjust the curve in-between those two points as you would normally. This approach uses the local adaptation method to compress the HDR image into something you can see, and then you use a custom tone curve to add some contrast back into the image. The results look similar to your best single-shot exposure (Figures 4 and 5), but they have additional detail in the shadow and highlights your single-shot couldn't capture.

After creating your compressed image, you will often want to apply some sharpening, and sometimes you'll also want to apply either a levels or curve adjustment to increase the contrast slightly.

Merge to HDR can help simplify a lot of complicated exposures, but there's one big problem. The source images all have to be from the same spot, and the subject can't move significantly. Since Photoshop is combining images to generate the HDR image, if you have clouds that move quickly or flowers that blow, they can affect the final image since there is no way to mask out part of the image before running Merge to HDR.

If the moving areas are small, you can use this technique to get a near-final image, and then, from your original image closest in exposure to the area of interest (such as the cloud that moved), clone the moving object into the final image. Given that the tool is brand-new in Photoshop CS 2, the people at Adobe are probably hard at work correcting this shortcoming for the next version!

Adobe has given us a great tool that lets you take pictures you would have shied away from before as being too hard to exposure, and it takes us one step closer to easily capturing what you saw in your mind and not just what your sensor read.



Figure 4: The best single-image exposure.



Figure 5: The combined exposure. Note the additional detail in the sky and forest.