Tonemapping HDR Images

Soft Tonemapping

With the HDR file open, the next thing we need to do is bring the highlights back into the range of what we can view on-screen. There are several ways to do this with the tonemapping tools in HDR Pro. Since we're going to do the rest of our tonemapping using the DZS, all we'll do in this case is use the Exposure and Gamma tonemapping adjustment and set the Exposure slider to 0. What you'll find in Photoshop when you open a 32-bit file and look at the Exposure and Gamma settings via View>32-bit Preview Options is that the Exposure slider is set to some value other than 0. What the value is will vary depending on the image. In this case it was +1.3. After zeroing out the exposure to bring the highlights back into visible and useable territory we have the image in figure 7.5, along with the histogram for the image.

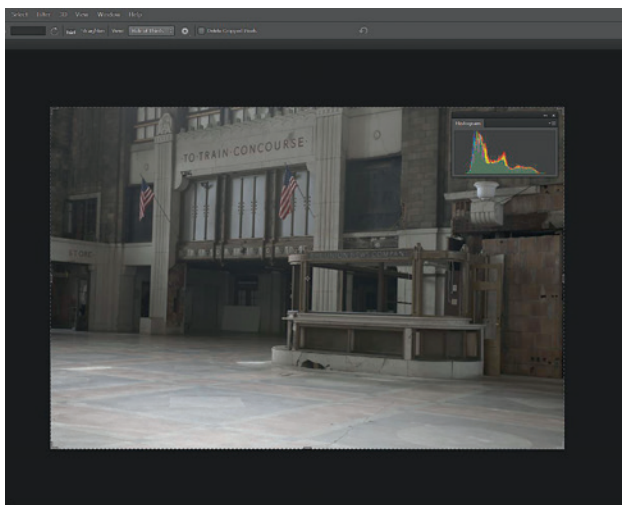

▶ *Figure 7.5: 32-bit HDR image with exposure set to 0 via Exposure and Gamma, and histogram*

You might look at the histogram for this image and the one for the single image and say the two are basically the same. They are similar, but there are important differences. The HDR histogram is more spread out and has a greater concentration of pixels in the darker areas. This is actually what we want because there is more shadow than highlight in the image. If you look carefully at the two versions of the shot, you can see that the floor area is a bit flatter and has better color in the HDR version. Overall, this is a better starting point.

I worked up two versions of this image, one using traditional tools and one using the DZS. The two are shown below for comparison (figures 7.6 and 7.7). The histograms for each are included as well. What do we see? The DZS version has better overall contrast and more visual "snap" than the traditionally edited version. There is better local contrast and global contrast. I've been better able to adjust the shadow and highlight detail. This can be seen clearly in the front of the newsstand. The DZS version has both deeper shadows in the wood and brighter light areas in the marble. I've also been able to adjust the

colors in the floor without affecting other areas because I could separate the adjustments based on luminance.

◀ *Figure 7.6: Image edited with traditional methods*

◀ *Figure 7.7: Image edited with the DZS*

Next, let's look at figure 7.8, which shows a screenshot of the layer stack for the DZS version. At the bottom, above the base layer, are the Hue/Saturation adjustments that were made to four selected zone masks. Above the luminance layer are the Curves adjustments that were applied. At the top is a slight global Curves adjustment to refine overall image contrast.

To see how I was able to isolate color according to luminance, let's take a closer look at the Red slider of the Hue/Saturation adjustment on Zone 7 of the DZS image (figure 7.9) and compare it to the Red slider of a global Hue/Saturation adjustment to the entire image when using traditional methods (figure 7.10). By pushing the slider all the way to 100 to overdo the effect, we can easily see which areas of the image are impacted.

We can clearly see that the adjustment to the reds is much more isolated to the floor in the DZS image compared to the other. Remember that we need light to see color. Colors in areas of shadow should be less vibrant than colors in areas of more light. That too is evident in the comparison of these two images. Notice the exposed brick on the left side of the archway to the train concourse. In the DZS image, the Red adjustment to Zone 7

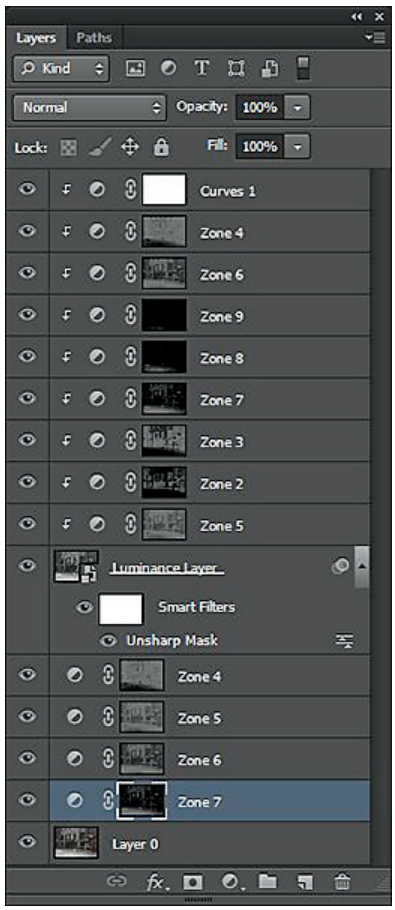

▲ *Figure 7.8: DZS layer stack*

doesn't affect the color in that exposed brick at all, whereas in the image using global Hue/Saturation adjustment the red brick is affected. We don't want that exposed red brick to be affected by the saturation adjustment to the floor because it sits in a shadow area and the color should be more muted. With the DZS and the ability to isolate adjustments based on brightness, we can achieve a more realistic and visually accurate result.

▲ *Figure 7.9: Red adjustment isolated to the floor in DZS*

▲ *Figure 7.10: Global Red adjustment*

Editing Images Post-Tonemapping

Can we use the DZS for post-tonemap adjustments to images that have been tonemapped in an HDR software application? Sure we can.

When I walked into the Central Terminal in Buffalo and saw the architecture, a concept came into my head of how I wanted to render the image. A visualization. What I wanted to do was to create something that evoked an old, Art Deco–era illustration. The image in figure 7.11 was my first attempt at realizing that concept, where I used the following traditional methods.

I loaded the source images into the HDR software program Photomatix from HDRSoft, tonemapped them, and saved the image as a TIFF, which I then brought into Photoshop for further editing. I applied Hue/Saturation and Curves adjustment layers, and sharpened the image with the High Pass method. I was pretty happy with the result. It was very close to what I was trying to achieve.

Next, I worked on the same image using the DZS. I used the same TIFF from Photomatix as the starting point. You can see the result in figure 7.12.

▲ *Figure 7.11: HDR image with global post-tonemap adjustments*

◀ *Figure 7.12: HDR image edited with the DZS*

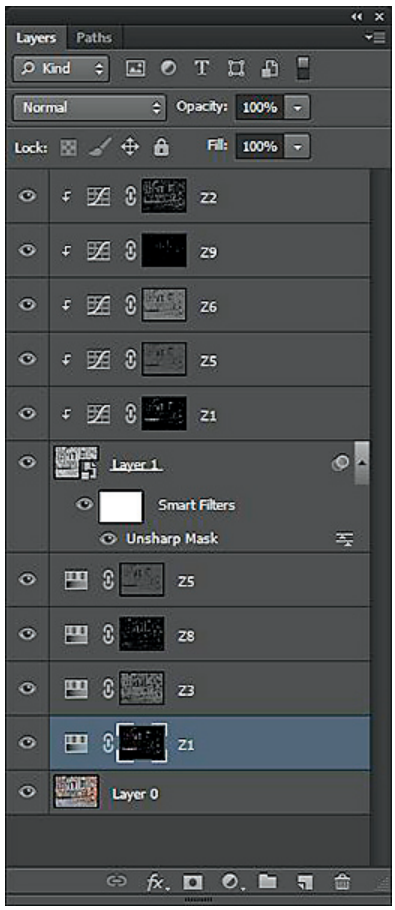

▲ *Figure 7.13: Layer stack for the DZS version*

This is the version I was looking for! In particular, the flags in the first attempt looked too realistic. In the DZS image I was able to give the flags more of an illustration look. Overall contrast is better and the image is sharper due to the fact that I was able to sharpen more aggressively on the luminance layer than I could when using the High Pass method on the traditionally edited version.

Some may look at this and say it just looks like an overprocessed HDR image, and that's OK. It is somewhat overprocessed compared to the strictly realistic interpretation we looked at earlier. It does, however, achieve the look I was going for and realize the concept I had in my mind when I walked into the space.

This is another key aspect of Adams' Zone System that can carry over to the DZS. As I discussed in chapter 1, visualization is important in being able to create a final version of an image that you're happy with. If you look at a subject or scene and can conceptualize how you want that image to look on-screen or in a print, you'll be much more able to edit the image effectively and efficiently. Visualization makes for a much more pleasant and productive editing session, as opposed to simply trying to take a shot and start making adjustments without a goal or concept. Work on visualizing the end result when you first view a scene, and the success rate of your images should improve immensely. The exercise in visualization will also help you in the field to compose the image with the final result in mind. This can help prevent having to do significant cropping after the fact to get a composition that is pleasing or to fit a composition to an otherwise directionless image.

Image Blending with Zone Masks

To finish off this chapter, we'll take a look at another way we can use our zone masks and zone groups. We can use them in blending exposures to extend dynamic range. There are several ways to manually blend images, but using the zone masks and zone groups can be very elegant.

Shot at sunset and into the setting sun, the image shown in figure 7.14 had a very wide brightness range. It was wider than I could capture in a single shot. If I exposed for the sky, I got an overly dark foreground. Exposing for the foreground produced an overexposed sky.

I needed some way to blend these two images to create a new version with proper exposure for both elements. I could have used HDR software. I could have used any of the other manual blending techniques as well, but those can be time consuming and inefficient due to the manual painting involved with some methods.

I decided to use zone masks to very quickly and effectively blend the two images together. I loaded them as layers with the foreground image (overexposed sky) on top. Pressing CTRL and clicking on the RGB composite channel in the Channels palette selected the highlights for the top layer.

As you can see in figure 7.15, it looks like a discrete selection. But remember that these are, in actuality, highly feathered areas. The marching ants only show where pixels of more than 50% visibility are selected. With the highlight mask selected, I simply hit the Delete key. The sky from the layer underneath shows through, and I have a blended image (figure 7.16) that is properly exposed for both the sky and the foreground.

Quick, easy, and effective. I could now go on and make further adjustments to either layer to refine the image more.

▲ *Figure 7.14: Exposed for sky (left) with dark foreground. Exposed for foreground (right) with overexposed sky.*

▲ *Figure 7.15: Highlights selected*

▲ *Figure 7.16: Blended image*

Chapter 8
Printing in the Digital Zone System

According to Ansel Adams' Zone System, printing is yet another area where you can exert control over the finished result. The goal of film exposure and development is to get what Adams called a "printable" negative. That did not necessarily mean a negative that would print well as a straight print. What it meant was a negative that contained the brightness range he wanted to reproduce in the print but left work to be done on the enlarger.

The printable negative Adams ended up with after making exposure and development decisions, if printed straight, would often render a flat, lifeless print. The reason is that the negative itself would have very low contrast. But Adams knew a few things about printing that led to the creation of that flat negative. He knew that he could do more work in the printing process to enhance the image on the negative and render a print with a broad tonal scale and good global as well as local contrast. He also knew that the dynamic range of the papers he used was narrower than that of the film, and thus he knew that if he developed a negative with the same broad tonal scale he wanted in the print, he would lose some of it in the printing process due to blocked-up shadows and blown-out highlights.

A well-known quote attributed to Adams is, "The negative is the score, and the print is the performance." Relating this quote to music, a songwriter may be able to write the music but it's the performance that brings the music to life. A straight, rote rendition of musical notes will sound flat, lifeless, and mechanical. For the music to have depth and soul it takes a strong performer to interpret the notes. In the same way, the flat, lifeless negative needs to be brought to life. On the enlarger, in the darkroom, Adams gave life to his negatives. By selecting a particular contrast grade of paper, by choosing the developer and dilution, and by using techniques such as dodging and burning, Adams was able to manipulate that printable negative and create a true work of art.

It's difficult to really understand the impact of Adams' quote without seeing some of his straight prints alongside some final versions of the same image. If you have the opportunity to do so, I would definitely recommend taking advantage of it.

With the DZS, you'll recall, I suggested that the two components are (1) exposure, and (2) development and printing. With digital, we don't have the opportunity to adjust anything at the time of printing. All the work needs to be done at the development stage. The print merely renders what we create in the digital darkroom.

The exposure decision gives us our RAW file. With that RAW file we work in the digital darkroom of ACR and Photoshop to create the image that will appear in the print. We do our color (or gray-tone) enhancement, and make global and local contrast adjustments to get the image ready to print. We do, however, still make those adjustments with the paper

choice in mind. With digital printing we don't have specific contrast grades of paper, but we do have media that can reproduce various levels of contrast and that have differing dynamic ranges and color gamuts.

The traditional and digital approaches are similar in that if either the negative or RAW file isn't right, you can't salvage something from it. You can't—either through development and printing of the negative or through editing of the digital RAW file—turn an unsuccessful image into a successful one. Remember the adage from earlier in the book—GIGO.

Choosing a Printing Surface

What I refer to as "hard" media are the coated papers. Glossy, semi-gloss, and luster papers fit into this category. They render images that are sharper, have greater Dmax (deeper blacks) and lower Dmin (brighter whites), and typically have a broader color gamut. Media that I call "soft" are papers like the watercolor or heavily textured, uncoated papers. Canvas also fits into this category, even glossy canvas. These papers don't produce as deep a black and typically not as bright a white, nor do they generally produce as wide a color gamut.

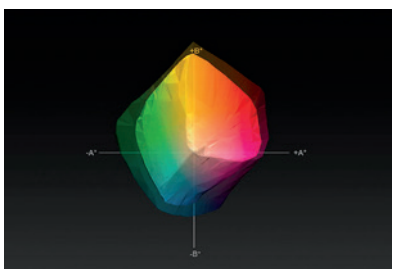

▲ *Figure 8.1: Color gamut comparison: Epson Exhibition Fiber (darker/larger plot) and Epson Velvet Fine Art (brighter/smaller plot*

Figure 8.1 illustrates a typical difference. In the illustration, the color gamut of Epson's Exhibition Fiber paper, a "hard" paper, is plotted against the smaller gamut of Epson's Velvet Fine Art paper, which is a very "soft" paper with a velvet-like feel. The difference between the two may not seem very great, but when you're printing an image with vibrant colors it is.

Knowing the characteristics of the paper allows us to choose the right media for the image and prepare the image so that it's reproduced as well as possible on that media.

The Right Paper for the Right Image

What do I mean by choosing the right media for the image? In my opinion, certain images work better on certain types of media. Crisp, clean, sharp images work better on hard media, whereas more muted or ethereal images with softer colors and less defined detail work better on soft media. An important aspect of choosing media is to avoid the cliché. There are fads in all areas of life, and certainly photography and art are no exceptions. For the past few years, canvas has been a popular choice for printing. It's relatively new to the photographer, and with digital printing it's readily accessible. Canvas can be used with wet prints, but that is a much more labor-intensive and error-prone process. Today, people are printing anything and everything on canvas because it's the hot trend. It's not because the image looks good on canvas—many images don't—but solely because canvas is popular. Metal is also becoming like canvas in that respect. Popularity isn't a good reason to choose a printing surface. Choose the surface that makes your image look the best.

Color Management for Printing

Recall in chapter 2, I said there was a color management aspect to printing and that we would cover it in this chapter. There is, and it's crucially important. Without a properly color-managed printing workflow, getting a print that comes out the way you want, with accurate color and contrast, is more good luck than good (color) management.

Just as with the monitor and scanner, we're going to make use of color profiles in the printing process. Without using these color profiles, paper coatings, optical brighteners that give a brighter white to the paper, any cool or warm tone the paper may have, and the way the printer interprets colors to combine inks on the paper, can all play a part in making the colors and contrast you see in a print different from what you see on-screen.

Most of the highly regarded paper manufacturers provide accurate profiles for their papers to be used with a variety of popular printers. Five or six years ago, these canned profiles were not as good, and having custom profiles made was often the best way to achieve high-quality results. Today, these manufacturer-provided profiles are quite good. Unless you're using a printer that the paper manufacturer doesn't provide profiles for, a paper from a lesser-known manufacturer, or your own papers, there really shouldn't be a need to have custom profiles made. That isn't to say manufacturers' profiles are perfect all the time. Sometimes there can be a dud. Recall from chapter 2 what you can do to try to rectify a problem with inaccurate prints. We'll talk about sending your prints out to a service bureau later in the chapter.

Soft Proofing

The first step in the printing process is something called soft proofing. Soft proofing allows you to get a good preview on your computer monitor of what your print will look like on paper. It doesn't necessarily replace a printed proof, and even if you do soft proof, your first print may not be exactly what you were aiming for. But the instances of reprinting should be reduced dramatically if you employ a sound soft proofing methodology.

The next section will walk you through the process of soft proofing in Photoshop. Later in the chapter we'll discuss soft proofing in Lightroom; the much-needed feature that was added in version 4. A word of advance warning, soft proofing in the two applications is a bit different so you'll want to pay particular attention to where you're working on your original and comparing to a copy (Photoshop) versus where you're working on a copy and comparing to your original (Lightroom).

▲ *Figure 8.2: Zoom Tool*

Soft Proofing in Photoshop

With your image open in Photoshop, go to Image>Duplicate. This will create a copy of your original that you will use to soft proof against. Tile the two images side-by-side, and make them the same magnification level. To change the magnification, click on the Zoom Tool (the magnifying glass) in the Tools palette (see the highlighted tool in figure 8.2). Moving your cursor over the image, you'll see the magnifying glass with a + sign in it. Click to increase magnification. Holding the ALT (option) key will turn the sign into a minus (−) and you can then click to reduce magnification. This function does nothing