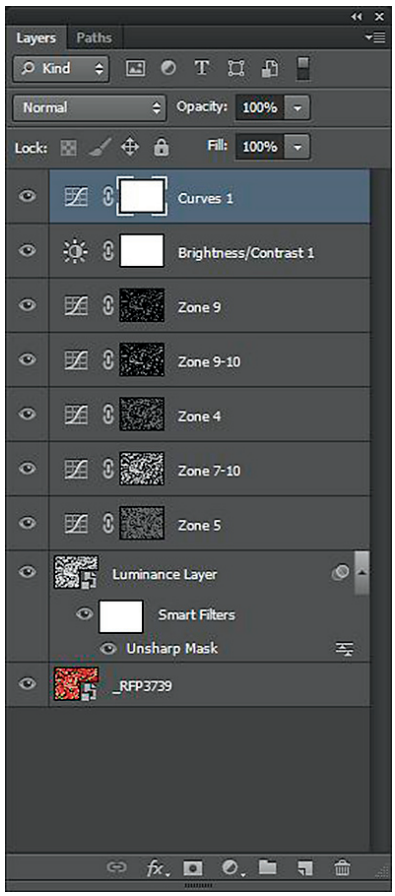

▲ *Figure 6.22: DZS pepper conversion layer stack*

bit of unimportant shadow detail at the bottom of the box to improve black levels. The layer stack in figure 6.22 shows what was done with this image.

In addition to the zone mask adjustments on individual zones to affect local contrast, I made some changes to larger tonal areas, as seen in the Zone 7–10 and Zone 9–10 adjustment layers. At the end, I made a global Brightness adjustment to bring down the entire image slightly, and then a global Curves adjustment to enhance contrast. Notice that these global adjustments are at the top of the layer stack. That's because I wanted truly global adjustments that took in the effect of the local adjustments.

ACR/Lightroom vs. DZS for Black-and-White Conversions

Let's take a look at one more color to black-and-white conversion. The image in figure 6.23 is the color version. It was taken in autumn on an overcast day. Because I used Daylight white balance, the bark has taken on a very interesting blue tone. It's almost otherworldly. Notice the difference in brightness, or luminance, at the front right of the image compared to the back left, where both areas have blue color. This will be important when we do our black-and-white conversions.

Working on the black-and-white conversions, I used the tools available in ACR (or Lightroom) on one and the DZS on the other. The ACR version is in figure 6.24 and the DZS version is in figure 6.25. The difference is not dramatic but it is evident.

In ACR we're trying to create gray contrast by working with color. We may be able to get the tonality we want in one part of the image but not in another, because adjusting, for example, the Blue balance will impact both the brighter areas at the front of the image and the darker areas at the back left. Instead, by extracting the luminance information from the image and using the zone masks to adjust brightness, we can more easily affect both local and global contrast in the image.

The DZS version took only a few minutes to complete. The ACR version took a fair bit longer because I had to go back and forth adjusting curves, Basic panel sliders, and color sliders to try to isolate the effects I wanted in specific areas. The work is definitely easier and more effective with ACR version 7 in Photoshop CS6 and with Lightroom 4 than earlier editions of the software programs because of the changes in the Basic panel adjustments discussed in chapter 4. Despite that, the DZS version has overall better tonality on both global and local levels as a result of the ability to isolate adjustments to specific areas of brightness.

Looking at the histograms for the two images in figure 6.26, we see noticeable differences. Compared to the histogram for the ACR version, the DZS version is more spread out, indicating a greater level of overall contrast. There are slightly more tones at the dark, left end of the graph as well as at the light, right end of the graph.

By now, what the Digital Zone System is and what it can do should be clear. In terms of converting a color image to black-and-white, the idea of creating good color contrast should translate into optimizing existing luminance contrast in a digital image. Extracting luminance and contrast that already exist in an image provides a better starting point for image editing, and thus leads to a better ending point, in most cases, than conventional methods.

▲ *Figure 6.23: Bark, color original*

▲ *Figure 6.24: ACR conversion*

▲ *Figure 6.25: DZS conversion*

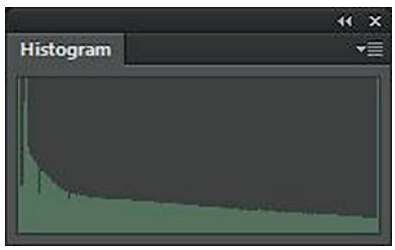

▲ *Figure 6.26a: Histogram for ACR conversion*

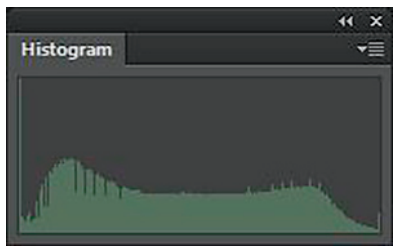

▲ *Figure 6.26b: Histogram for DZS conversion*

Chapter 7
The Digital Zone System and High Dynamic Range Imagery

- Ansel Adams' Zone System was analog High Dynamic Range imagery (HDR).

- HDR is an extreme version of ETTR.

Controversial statements? We'll see. Over the course of this chapter we'll look at why I believe those statements to be true and at how the DZS can be incorporated into an HDR workflow. Perhaps by the end of the chapter you'll agree that these two statements are true as well.

Adams' Zone System is a method for controlling dynamic range and contrast in negatives and prints. Because he understood the medium he was working with and the characteristics of each film, he could carefully adjust exposure and development of the film to produce a "printable" negative. Why have I put printable in quotations marks? We'll examine that more in chapter 8.

With the adjustment of exposure and development, Adams could compress the brightness range of a scene into something that would fit within the range he could capture on the film. Conversely, he could expand the brightness range of the scene on the film.

The option always exists to expose normally (i.e., according to the meter reading) and develop normally. But with the Zone System there are a couple of other options as well.

We can reduce exposure and extend development. We can also increase exposure and reduce development. Whichever we choose depends on the brightness range of the scene, which Zone we place important shadow detail in, and how that decision causes highlights to fall relative to the brightness range of the film.

There is no free lunch, however. You can't cram a wider brightness range onto the negative than it can handle. Just as with RAW capture in digital, there is room at the highlight end of the scale, and where placement of a shadow detail might cause some highlights to fall into a higher zone, they can be brought back into range within reason. This is the essence of ETTR. On the other end, you can't extend a midtone in a low-dynamic-range scene into Zone VIII or Zone IX. You may be able to lift it to Zone VI or VII, but that's about it.

In the context of HDR, we're going to concentrate on compression of dynamic range. Even with the more limited ability to compress dynamic range in the Zone System, what is happening is tonemapping for all intents and purposes.

If you read Adams' writings on the Zone System, particularly The Negative, he uses terminology that is similar to what we use in talking about HDR tonemapping. Adams discusses global and local contrast, both of which are key components of editing or

tonemapping HDR images. He discusses the impact local contrast has on the feel of the image. Lower contrast, on both a global and a local level, produces an image that is flatter and has less texture. Higher contrast, particularly at the local level, enhances textures and can add a more gritty look to an image. This is precisely what happens when we tonemap an HDR image.

There are other methods for combining two or more images to increase dynamic range. Manual blending of images has been done since the mid-1800s when Gustave Le Gray combined negatives at different exposure settings to expand the dynamic range of his images. Manual blending of images in Photoshop is also possible using layer masks and selectively revealing certain parts of images. In this chapter we're going to concentrate on HDR as an image-blending and dynamic range expansion method. Later in the chapter there's a section on using your luminance masks to blend images as well. One of the knocks against HDR is that it can't produce "realistic" results. This is, of course, nonsense, and some of the examples in this chapter will show that realistic results can be achieved.

Some will say there is no need for any sort of zone system, digital or otherwise, because we have HDR and other advanced image-blending techniques. The thinking goes that since the purpose of the Zone System was to manipulate dynamic range, HDR and similar methods serve that purpose. The argument continues that there is no need to really understand metering because you can just blast away as many brackets as you want and choose later which ones to merge. That really misses the point of what HDR is about. It also shows a clear lack of understanding of the Zone System. In addition, it relies on a method of photography called "spray 'n' pray," in which people just crack off rapid-fire image sequences using the high burst rate and buffers of digital cameras in the hope that they get a useable image, or useable set of images, out of it. That's not the approach that a serious photographer wants to take in making thoughtful and deliberate images.

Many people do their HDR merge, tonemap it into something surreal looking, and call it a day. Others take a different approach, tonemapping in the HDR application to a more natural look but stopping the process at that point. There is no rule, however, that says you have to tonemap within the HDR application or use the tools the HDR software provides. There is another way to look at that HDR file. That is to consider the HDR file as a new RAW image. That's right; I'm saying HDR is the new RAW. There are times when I will do everything in the HDR application. I'll merge, tonemap, output a final image, and be done. Those instances are rare. Very rare. Those instances will typically be in situations where I want to get a quick preview of what I may end up with, even when I'm trying for a more surreal look. What I'm more likely to do is merge; tonemap the 32-bit image to a fairly flat, low-contrast result that gives me plenty of options for adjusting contrast later; then work on creating a final version that has the full range of contrast I want it to have. This flat, low-contrast "working image" is akin to Adams' "printable" negative. Sometimes I won't even tonemap in the HDR application at all. I will simply merge, open the 32-bit image in Photoshop, drop it down to 16-bit without any HDR tonemapping, and undertake a process called soft tonemapping, using the tools available in Photoshop with the DZS.

Another way to look at an HDR file is to consider it a new RAW image.

How do we work with RAW files currently? For the most part, we make adjustments to the file in our RAW converter of choice, then open it in our image editor, make additional adjustments with layers and layer masks, and save that final edited version as our master in either TIFF or PSD format. Some work exclusively in a hybrid program like Lightroom and never touch Photoshop. Lightroom 4.1 can work with 32 bit files in the TIFF format but you still can't work with 32 bit .hdr or .exr file. For all intents and purposes, Lightroom really still can't be considered to be compatible with 32 bit HDR images.

What we will work through is a way to use the DZS as a method of tonemapping HDR images. The biggest upside to this approach is that you don't run the risk of getting the overprocessed, surreal results that are possible with regular HDR tonemapping tools. The results are completely natural—which is a look a lot of people want but have difficulty getting with the regular HDR tonemapping methods. We'll also look at using the DZS for editing images that have already been tonemapped with standard HDR tools.

This chapter is not going to be an in-depth treatise on HDR. If you would like to delve into HDR in more detail, then *Practical HDRI, 2nd Edition: High Dynamic Range Imaging Using Photoshop CS5 and Other Tools* by Jack Howard is recommended reading. What we'll discuss in this chapter is how the DZS can fit into an HDR workflow.

Employing the DZS as a method of tonemapping is simple and requires only a few additional steps compared to working with it on non-HDR images.

What is HDR?

Before getting too far along, a brief explanation of what HDR is is probably in order. HDR in a purely technical sense is the merging of a series of images of differing exposures (differing brightness values) into a 32-bit image file. The goal is to extend the dynamic range of the image beyond what the sensor can record in a single exposure. Manually blending images is often referred to as HDR, and it is in the sense that manual blending also extends the dynamic range. But it doesn't involve the use of a 32-bit file format, so from an absolute technical standpoint it is not HDR. A 32-bit file format is used because it can contain a much wider range of brightness than a16-bit or 8-bit file.

We can't currently use these 32-bit HDR files. Existing monitors can't display the entire brightness range in 32-bit images, and printers can't reproduce the entire brightness range with existing technology. Perhaps one day we'll be able to use HDR images, but not today. In order to make use of the files, we have to take them through a process called tonemapping. Tonemapping is just a fancy word for editing. In this process the wide brightness range in the HDR image is compressed into a narrower band that we can actually use. This "mapping" of tones is in essence no different from what Adams did with his exposure and development decisions. With that short explanation of HDR, let's move on to seeing how it fits within the DZS.

The goal of HDR imaging is to extend the dynamic range of an image beyond what the sensor can record in a single exposure.

HDR Merging

We'll start by looking at the RAW file we get out of the camera. Whether you're using the HDR Pro utility in Photoshop or third-party HDR software, there are a few things you should do to RAW files before sending them into HDR-land. Yes, I am recommending making some edits to the RAW files before doing the HDR merge. I know others will say that you should do nothing to the RAW files, but I disagree. In my opinion, there are some adjustments that are simply better to do at the RAW stage.

The reason I suggest making these adjustments at the RAW stage is due to the nature of the HDR process. Everything gets enhanced, whether good or bad. If there's something not quite right in the RAW images, it will get exaggerated in the HDR merge and may be more difficult to correct later.

The first thing to correct is the white balance. If the white balance setting in the camera was not set properly and your image has an unwanted color cast, fix it at the RAW stage. In either Lightroom or ACR you can adjust the white balance on one image, and then sync the changes to the other images in the sequence. This is probably the most important pre-merge adjustment of the three I'm suggesting. If you end up with an unwanted color cast in the HDR image due to an improper white balance, it can be difficult to deal with.

The second thing to take care of before merging are aberrations like purple fringing and chromatic aberration. It's quicker and easier to fix them in the RAW files, and the result is better than if you try to fix them later. Since we're working on multiple copies of the image in our layer stack in the DZS, trying to fix aberrations later can cause some unpleasant results due to the way pixels are shifted to correct the errors. Again, you can do it on one image, then sync to the rest.

The last pre-merge adjustment is spotting (to invoke a film term) or cloning/healing of dust and other detritus in the source images. Of the three this is probably the least necessary, but I find it's still better done at the RAW stage than later in the workflow. As with the other two flaws, you can fix dust problems in one image, then sync to the rest in the series.

Once you've done your pre-merge edits, you can go ahead and create your 32-bit HDR image. In many of the examples that follow, the 32-bit images were created in HDR Pro. You can use any HDR software for this. I like HDR Pro because of the seamless integration with Lightroom, Adobe Bridge, and the DZS.

RAW vs. TIFF for HDR Processing

Whether you send RAW files or TIFF files to your HDR application of choice really depends on what the particular application works with better. Most HDR applications have built-in RAW converters, but the quality of the conversion can be variable. For the best results, I recommend sending the TIFFs you create from your edited RAW files in 16-bit, tagged with the ProPhoto RGB or Adobe RGB color space.

If you use a program like Photomatix and its Lightroom plug-in, when you export the files to Photomatix from Lightroom they will be converted to TIFFs or JPEGs, in 16-bit or 8-bit, and they'll apply whatever color space you selected, along with the edits you made to the RAW files. HDR Pro in Photoshop will automatically pick up edits you made to RAW files by reading the XMP sidecar files.

Next, we'll walk through a series of images that take you from the original 32-bit HDR to a final image, with a comparison between the DZS version and a traditionally edited version.

The first example is an image of an old, disused railway station in Buffalo, New York. The Buffalo Central Terminal is a terrific example of Art Deco design, and the station was an important hub for both passengers and cargo in its day. Unused for many years, it has fallen victim to heavy vandalism and is in the process of being restored by a local preservation group.

In figure 7.1, you can see what the 32-bit version looks like before any adjustments. The histogram for the image is included as well.

◀ *Figure 7.1: Original 32-bit image and histogram*

What should be immediately clear from the shape of the histogram is that it very much resembles the histogram of an image that has been exposed using ETTR. Remember that earlier in the chapter I suggested the Zone System was an analog version of HDR and that increased exposure and reduced development represented an early form of ETTR? In this histogram we see that quite well. Remember that I also said HDR was an extreme version of ETTR? This histogram illustrates that too.

The dynamic range of this scene was not very large—only about 11 stops or so. In a purely technical sense, I probably didn't need to use HDR because that range fits within what the sensor in the camera I was using (a Nikon D700) can capture. But what gets tested in a lab and what happens in real-world shooting conditions can be different, so I wanted to take the precaution of giving myself the leeway of those extra exposures and the ability to merge to HDR. Due to the many windows in the space and the type of day it was, with rapidly moving clouds, the light in the building was changing often and quickly. What might have been just an 11-stop range one minute could have been a 13-stop range or more the next. Before moving on, let's take a look at a single exposure to see what we have.

Single File and Adobe Conversion of RAW Files

A screen capture of the middle exposure from the bracket (the "normal" exposure) in ACR is shown on the left in figure 7.2. We can see that even this normal exposure is pushed to the right. The reason is that the area I was metering was fairly dark and my meter wanted to turn that dark area into a middle-toned value. There is no highlight clipping, so we can still use this image.

As discussed in chapter 4, with Lightroom 4 and ACR 7, Adobe introduced a new RAW conversion engine. After some adjustments with the older RAW conversion engine we get the image in figure 7.3. It's better. The histogram has been shifted so it's just off the right edge. While there is still good shadow detail, the floor is overly bright. This is about the best I can do with the old process.

Switching to the new PV2012, we get the image shown in figure 7.4. There is now much better detail in the floor and the shadow detail is still preserved. Looking at the histogram, you can see it has very flat contrast overall. I could work with this version and use the DZS on it. But it's still not as good a starting point as the HDR version we'll look at next.

◀ *Figure 7.2: Unaltered RAW image from the camera*

◀ *Figure 7.3: RAW image edited with Process Version 2010*

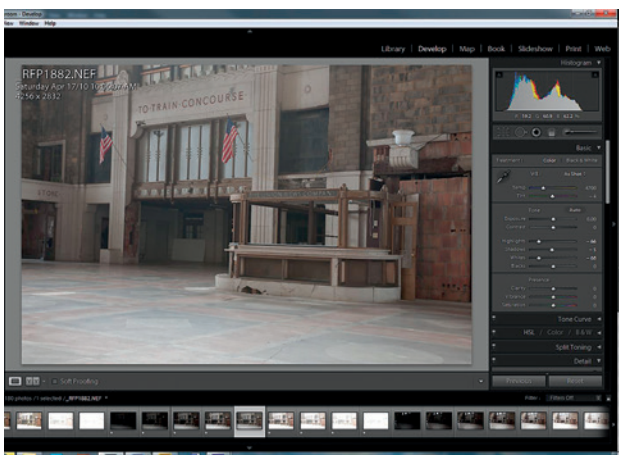

◀ *Figure 7.4: RAW image adjusted with Process Version 2012*