

# Creating HDRI Panoramas

*There are different methods using different software to create a spherical HDRI panorama. I've tested several methods and scrutinised the results. Here is an overview and the workflow that appears to work best for me.*

## Introduction

I have gathered some experience in HDRI panorama capture over the years, changed equipment and experimented with different software tools — and read a couple of books on the subject to get some insight into the topic. Though this is centred on my equipment, most of what transpires here can be adapted to any equipment.

## Camera, Lens and Tripod

Using the Nikon D50 6 Megapixel DX camera with the Nikon 10.5 mm 1:2.8 G ED AF DX Fisheye gives nice and fast results up to 8,000 by 4,000 pixels in the spherical projection with 8 shooting positions. The AF-S Nikkor 10-24 mm 1:3.5-4.5 G ED DX lens can be used at 24 mm for panoramas up to 18,000 by 9,000 pixels with 42 shooting positions.

Using the Nikon D600 24 Megapixel FX camera with the Nikon 16 mm 1:2.8 D AF FX Fisheye gives nice and fast results up to 16,000 by 8,000 pixels in the spherical projection with 8 shooting positions. The Nikon 8-15 mm 1:3.5-4.5E ED AFS FX Fisheye gives a bit smaller panoramas but needs less exposures. The AF-S Nikkor 18-35 mm 1:3.5-4.5 G FX lens can be used at 22 mm for panoramas up to 21,000 by 10,500 pixels with 16 shooting positions.

The Manfrotto 190XPROB tripod with the 438  $\pm 10^\circ$  ball head and the spherical VR head 303SPH proved to be a stable and reliable tool to hold the heavy camera and lens.

## Software for File Conversion, HDRI Merging and Panorama Stitching

The sensor of the Nikon D600 covers 14.3 EV and those are saved in a RAW file. Saving as JPG, only 8.5 EV remain — 60% of what the sensor could see. The RAW files (NEF: Nikon Electronic Format, a compressed TIFF variant with a couple of proprietary tags) can be read by the Nikon View NX 2 program that comes free with the camera. With this tool, the NEF files can be converted to 16-bit TIFFs.

Adobe Photoshop CS5 with the help of the Bridge can merge exposure series and aligns them automatically. Optionally, they can be de-ghosted. Pictrenaut 3.0/3.2 can also merge exposures to an HDRI, aligning and de-ghosting can be enabled. HDRShop 1.0.3 can be used to merge exposures to an HDRI as well. I have also used Artizen from version 2 on, meanwhile Fhotoroom 3.0 and Luminance HDR 2.3 (formerly qtpfsgui).

As for stitching the individual shots to a spherical panorama, there are Hugin and PTGui Pro that I have used and PanoStudio2 Pro for cylindrical panoramas. There are others like Auto-Pano and Pano Tools with which I have barely experimented.

Panorama projection transformers I know are HDRShop, PanoCUBE, GoCubic, Cube2Cross, Pano2QTVR, PanoVR and the Photoshop Flexify2 plug-in.

## Literature

I have read a couple of books on the subject that helped me hugely understanding the subject of HDRI and IBL. Some are in German, some in English. A list is included at the end of this document.

## Camera Setup

There are settings defined in the camera menu and others that have to be varied when acquiring HDRI images. The important settings are:

### Semi-permanent settings in the Camera Nikon D600

Quality:	RAW
Size:	full (6016 x 4016 pixels)
Colour depth:	14-bit
White balance:	direct sunlight
Auto distortion correction:	OFF
Colour space:	Adobe RGB
Vignetting correction:	normal
Noise reduction:	ON for exposition times of 1 second and longer
Noise reduction ISO+:	normal
ISO sensitivity:	100
ISO automatic:	OFF
Remote control delay:	2 seconds
Timed shutter release:	6 seconds
Exposition increments:	1/3
Bracketing:	only exposure time (for experimenting)
Bracketing sequence:	under – correct – over exposed
Bracketing distance:	2.0 EV

### Ad-hoc settings for exposure series

Light conditions change and need to be addressed on the spot. Outdoor panoramas with a visible sun need more exposures than if the sun is behind a building or a tree trunk. Cheating the sun in afterwards is an option I do not consider. Indoor scenes are yet different.

The D600 does feature bracketed exposures but unfortunately only three. It is easier not to use the automatic bracketing option. After experiments, I abandoned the use of this option.

## Lens Setup

Make sure the lens is clean and, more importantly, there are no dirt specks on the Bayer filter in the camera. Those are mostly responsible for the blotches. The camera settings included the f/stop, which is a quality of the lens, not the camera. However, many lenses have no manual f/stop ring anymore and are controlled by the camera. This is quite convenient.

All lenses I use show the best sharpness, low chromatic aberration and acceptable vignetting at f/8 – vignetting is corrected by the stitching program anyway. The Nikon 16 mm 1:2.8 D AF FX Fisheye must be locked with a slider at f/2.8 so that the camera can control aperture.

DOF (depth of field) is a moderate issue. For both fisheye lenses, the 10.5 mm DX and the 16 mm FX, setting the focus distance to 1 m results in a DOF at f/8 of around 55 cm to 21 m, which is quite suitable of an indoor panorama. For outdoors, set the focus to infinite to get a DOF range from 1 m to infinity. On both lenses infinite is a hard stop. The Nikon 8-15 mm 1:3.5-4.5E ED AFS FX Fisheye is very difficult to focus manually. Pointing to a faraway object sets it to infinite and it is sharp from 1 m to infinite at f/8 (checked at 14 mm). The lens hood must be removed when a focal length below 14 mm is used.

For the 18-35 mm FX lens, the scale is quite accurate from the shortest distance to the 1 m mark. The distance should be set between the 1 m and the beginning of the infinite symbol mark for an indoor panorama. For outdoor ones set distance to the thickest part of the first part of the infinite symbol. Always start adjusting from near to far because if you start at infinite distance, the scale is not the same. Additionally, the focus breathes, changing the focal length a bit. This lens can be set to a focal length of 18 mm easily because it stops at its shortest length. 22 mm is the maximum that can be used to shoot a panorama with only 16 positions. This is the extruding rib exactly between the 18 and 24 mm marks. I have not yet used this lens with a longer focal length for panoramas.

### Tripod and Panorama Head Setup

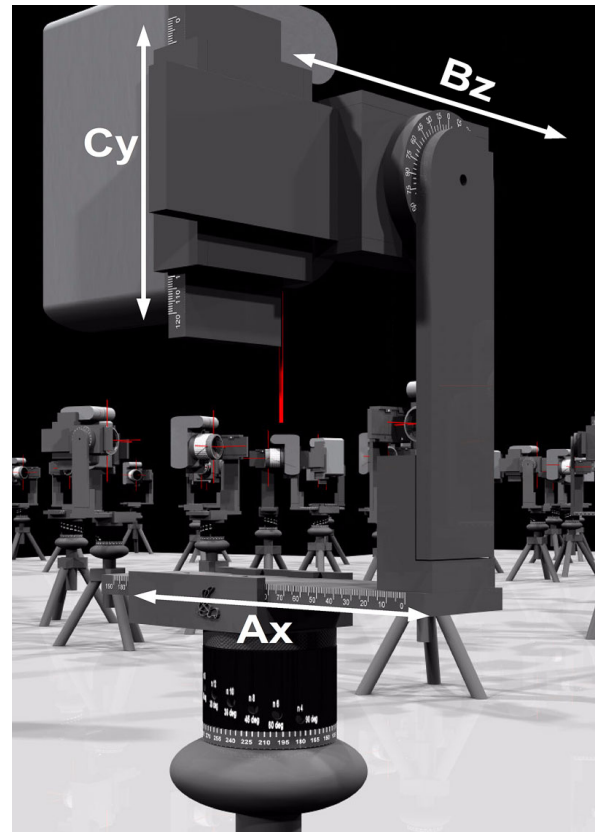
Depending on the lens used to shoot a spherical panorama, the panorama head has to be adjusted in all three axes X, Y and Z in order to get a panorama that can be successfully stitched. There are no shortcuts and the settings on the panorama head have to be checked against the table below.

Adjusting the panorama head for each lens for the first time is quite time consuming but once the positions are established, they remain the same for the same lens-camera assembly.

The picture at right identifies the axes. All three axes have an mm scale engraved, making it easy to adjust with a precision of half an mm.

- **Ax** is the fundamental horizontal slider that is fixed on the horizontal rotating head and it shifts the camera left and right in the X-axis. It has to be adjusted according to the camera used.
- **Bz** is the slider that moves front and back, the Z-axes. It has to be adjusted according to the camera used. It must also be adjusted according to the lens used and if it is a zoom lens, according to the focal length.
- **Cy** is the vertical slider on which the camera is mounted. In fact, this slider is mounted on the camera itself. It has to be adjusted according to the camera used.

It should be possible to keep the settings of Ax and Cy if the same camera is used. However, it depends on the precision of the lenses used.



### Lenses FX for the D600 Camera (Block Z-fix 181 mm = 18 mm to end)

Lens	Ax (mm)	Bz (mm)	Cy (mm)	Steps (°)	Row 1 (°)	Row 2 (°)	Row 3 (°)
FX 8-15 mm FE 8 mm	99.0	15.0	32.5	120	0 x 3	+90 x 1	-90 x 1 (+1)
FX 8-15 mm FE 12 mm	99.0	15.0	32.5	90	0 x 4	+90 x 1	-90 x 1 (+1)
FX 8-15 mm FE 14 mm	99.0	15.0	32.5	72	0 x 5	+90 x 1	-90 x 1 (+1)
FX 8-15 mm FE 15 mm	99.0	15.0	32.5	60	0 x 6	+90 x 2	-90 x 2 (+1)
FX 16 mm Fisheye	99.0	42.0	32.5	60/0	0 x 6	+90 x 2	-90 x 2 (+1)

### Creating HDRI Panoramas

FX 18-35 mm, 18 mm	99.0	12.0	32.5	45/45	-30 x 8	+30 x 8	+/-90 x 1 (+1)
FX 18-35 mm, 22 mm	99.0	13.0	32.5	45/45	-30 x 8	+30 x 8	+/-90 x 1 (+1)

#### **FX Panorama Sizes** (without nadir)

FX 8-15 mm Fisheye at 8 mm: 4 shots, 7,500+ x 3,750+ pixels

FX 8-15 mm Fisheye at 12 mm: 5 shots, 11,800+ x 5,900+ pixels.

FX 8-15 mm Fisheye at 14 mm: 6 shots, 13,600+ x 6,800+ pixels

FX 8-15 mm Fisheye at 15 mm: 8 shots, 15,000+ x 7,500+ pixels.

FX 16 mm Fisheye: 8 shots, 16,000+ x 8,000+ pixels.

FX 18-35 mm at 18 mm: 17 shots, 19,400+ x 9,700+ pixels.

FX 18-35 mm at 22 mm: 17 shots, 21,600+ x 10,800+ pixels.

#### **Lenses DX with the D50 Camera** (Block Z-fix 171 mm = 27 mm to end)

Lens	Ax (mm)	Bz (mm)	Cy (mm)	Steps (°)	Row 1 (°)	Row 2 (°)	Row 3 (°)
DX 10.5 mm Fisheye	99.5	44.0	32.0	60/0	0 x 6	+90 x 2	
DX 10-24 mm, 10 mm	98.5	19.0	34.5	45/45	-20 x 8	+50 x 8	
DX 10-24 mm, 12 mm	98.5	20.0	34.5	45/90	-27 x 8	+45 x 8	+90 x 1
DX 10-24 mm, 15 mm	98.5	20.0	34.5	36/72	-33 x 10	+33 x 10	+70 x 5
DX 10-24 mm, 18 mm	98.5	18.0	34.5	30/60	-39 x 12	+15 x 12	+60 x 6
DX 10-24 mm, 24 mm	98.5	10.0	34.5	30/60	-45 x 12	0 x 12	+48 x 12 +60 x 6
DX 18-55 mm, 18 mm	100.0	23.0	33.0	36/60	-40 x 10	+15 x 10	+60 x 6

**Note:** DX 10-24 mm at 18 mm is better quality than DX 18-55 mm at 18 mm. The 16 mm FX Fisheye is way better than the 10.5 mm DX one even with a 0.6 ND gel filter on the rear. The 8-15 mm FX Fisheye is yet better than the 16 mm Fisheye.

#### **DX Panorama Sizes** (without nadir)

DX 10.5 mm Fisheye: 8 shots, 8,200+ x 4,100+ pixels.

DX 10-24 mm at 10 mm: 16 shots, 8,200+ x 4,100+ pixels.

DX 10-24 mm at 12 mm: 17 shots, 9,600+ x 4,800+ pixels.

DX 10-24 mm at 15 mm: 25 shots, 12,200+ x 6,100+ pixels.

DX 10-24 mm at 18 mm: 30 shots, 14,500+ x 7,250+ pixels.

DX 10-24 mm at 24 mm: 42 shots, 18,800+ x 9,400+ pixels.

#### **Positioning the Tripod**

Because the tripod obscures the nadir it should be set to a location that makes it easy to cheat the nadir into the panorama afterwards. The Manfrotto panorama head is great but obscures the nadir generously. The more level it is, the better the panorama can be stitched together.

First, look from where the brightest light shines and what shadows the tripod casts on the ground. Rotate it so that the shadows of two legs fall together so that the shadow of only two and not three legs must be removed.

On an uneven ground, adjust the height of the three legs to get the head as level as possible. Check the bubble level. Then adjust the ball head, also with the help of that bubble level. Mount the camera with the lens. Check and if necessary adjust Ax and Cy, adjust Bz in any case. Make sure the tripod cannot slide or tilt in any way; push the tripod down into a soft

ground. Do not extend the centre part if it can be avoided. The top of the tripod with panorama head, camera and lens is quite heavy and tends to swing.

Remove the cover on the flash contact and put yet another bubble level on it. On the panorama head, set pitch to 0° so that the lens looks horizontal and fine adjust the ball head. Rotate the camera on the head (Yaw) and make sure the bubble level is level when the camera looks in any of the four cardinal directions North, East, South and West.

Ascertain that the locking screw is in the right slot for the rotation interval. Usually at 60° for the 16 mm fisheyes and 45° for the rectangular lenses mentioned. For 72°, use 36° and for 120° use 60° and move two stops. There is a 90° stop for the 8-15 mm Fisheye at 12 mm.

## **Shooting the HDRI Panorama**

The tripod is placed and level, the camera with the clean lens is mounted and adjusted correctly on the panorama head. It is time to get it over with.

### **Where to put the transportation gear**

It happens that you discover the camera bag, the backpack or whatever in the panorama. If you put anything down on the ground — which is a very bad habit — move it each time the camera is rotated. This is something that gets usually forgotten; been there done that. Better, transport the equipment in a backpack. Take it out and set the gear up, then pack everything not needed for the panorama shooting into the rucksack and put it on your back.

Be careful when going into hiding that your backpack does not sneak out behind the foliage of the bush behind which you are hiding. I discovered such once to my horror but it was luckily on the edge of the shot and disappeared in the stitched panorama.

### **Stability and Focus**

Always check that all screws are tightened after the panorama head was moved. Make sure the distance on the lens is set correctly. You may also focus automatically to the nearest and the farthest objects and put the distance in between, but this is usually not better than taking the values mentioned above. Do not forget to disable auto-focus on the camera and the lens. Never touch the lens again until the shooting session is over.

### **Vibrations**

When shooting watch the bubble level on the flash contact and start exposing only if the bubble is rock stable. Indoors, check how much the camera vibrates on the tripod when you move around and behave accordingly.

Vibrations may also be due to a heavy lorry or a train passing nearby, even if they are not in the shot. If there is an earthquake, you have probably other worries than keep on shooting the panorama; besides, you may have to start the sequence anew if the environment changed.

### **Moving targets — potential ghosts**

Watch out for cars, pedestrians, animals, birds, clouds, foliage moving in the breeze. Make sure the sun will not be covered by a cloud in the next 10 minutes or so. Breeze is very annoying. There is none where you stand and the vegetation nearby is stable as a rock, but a couple of ten metres away or the crowns of trees still catch a breeze and you end up with double or triple twigs. And if there are clouds in the sky, those buggers are also pushed by a wind so high up that you do not feel it.

## Moving light and shadows

The sun moves 15° per hour, this is a fourth of a degree per minute or 2.5° in ten minutes — and so do the shadows. There is this dilemma: shoot fast to keep sun and shadows mostly at the same location but take your time to watch out for things that can get into the shot, even far away, and wait until the breeze dies out for a moment. Indoors are much less problematic in this respect provided you can keep visitors out of range.

## Shooting sequence

A proven method is to align the main light source (e.g. the sun) in the centre of the image, then tilt the camera 90° up and take the zenith shots. Check with the bubble level on the camera or use the virtual horizon of the D600 (near the end of the Systems menu).

Then move the camera back horizontal and take the 0° brackets. Move the camera one step to the right (eastwards, clockwise) and shoot the next bracket series. Repeat until the circle is full. Point the camera down and make one nadir bracket, mostly the bulky tripod. Rotate the camera around about 90° and redo another nadir bracket series. If a rectangular lens is used, shoot the lower row first, then the upper row.

If the tripod is on a patterned floor that cannot be cheated in later, move the tripod away by about 120 cm, 1½ time of the leg distance, tilt the camera down about 45° and make a bracket series with the spot where the tripod was approximately in the centre. Watch out for reflections on the ground that was not there when the tripod was there.

## Shoot more than ultimately needed

When in doubt a setting was really shot or only believed to have been shot, shoot the whole bracket series again. It is less tedious to identify a full bracket sequence than finding where a shot was missing and added. If in still greater doubt, redo the whole panorama. And after that long hike to the spot, where is the need to get hasty all of a sudden?

Experience showed that the auto-bracketing does not always give the expected results. Adjusting the exposition time for each shot and trigger with the remote control is reliable.

## Developing the photographs

The photographs are digital, not film, but they need to be developed or prepared nevertheless. Particularly the NEF files from the D600 cannot be read by most of the programs and those that can only interpret the image data but not all the camera data embedded in undocumented proprietary tags. There is no better way but using a Nikon proprietary program like the free View NX 2 or the rather expensive Capture NX 2.

## Create a folder structure for the files

There will be a lot of files messing up the hard disk and confuses you in no time. Create a folder structure to organise all files throughout the project. I created a folder structure template I can copy it into the new project. I settled to this structure:

*/project name*

*/HDRI*

Stitched HDRI panorama and temporary cube faces and processed variants

*/Movies*

Finished QTVR, Flash and HTML5/Javascript movies

*/Originals*

## Creating HDRI Panoramas

NEFs, there might be a subfolder for the pictures that are redundant and deleted later

*/Pict*

Images and thumbnails for the web pages and other entertainments

*/Probes*

Light probes in different sizes (and projections if appropriate)

*/StitchedPano*

LDRI raw panorama, files generated by the stitcher, reference curves, etc.

*/TIFF (deleted once the project is done)*

Converted NEFs and merged individual HDRIs

*/Tonemapped*

Tone-mapped variants

A text file with the technical details for each picture making up the HDRI panorama

Another text file with the data that are sold along with the light probes

### Getting the exposures from the camera

Just get the Ultra SD HC1 32 GB SanDisk memory card out of the camera, put it into a slot in the computer and copy the files onto the hard disk (HD) into the */project/Originals* folder.

### Camera and lens data

Nikon's programs are a bit clumsy to use but View NX 2 can do the job nicely. Open the folder with the downloaded exposures, they get displayed as thumbnails. Find any rogue images and move them to another folder so that only the valid photographs are displayed. Select them all by clicking on the first one, hold the [Shift] key and click on the last.

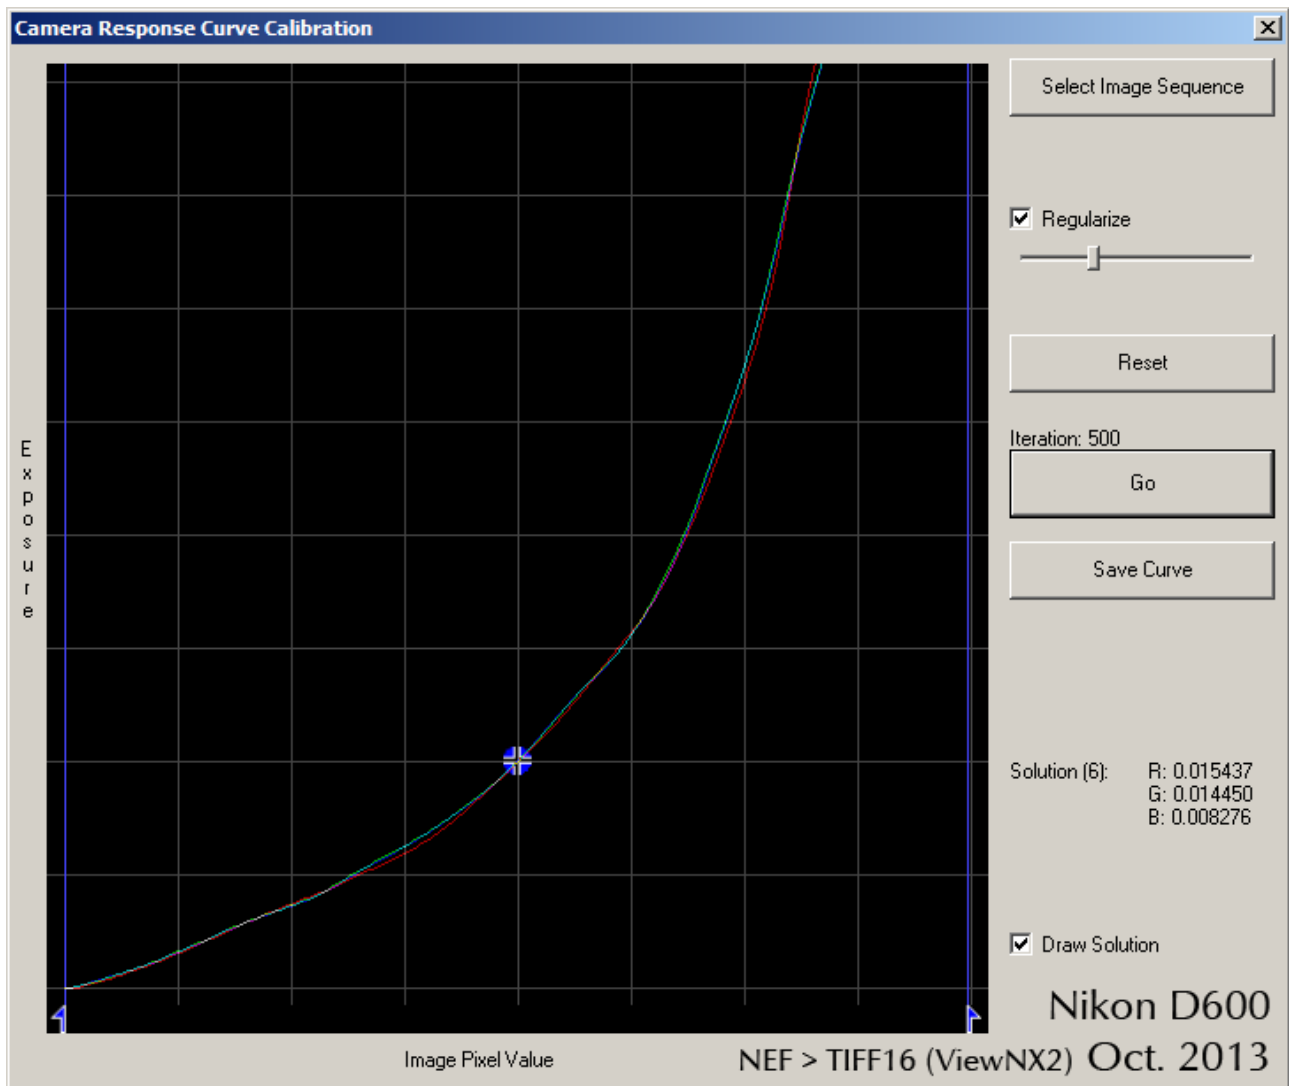
Either click on the right side to open the Bearbeiten sliding window or open it via the menu *Fenster > Bearbeitung*. Select the tab *Anpassungen, alle Anpassungen* [sic!] and scroll down to the end. Make sure *Farblängenfehler* and *Farbquerfehler* are ticked: this is the horizontal and vertical colour aberration correction.

Convert all NEFs to 16-bit TIFFs. There are three ways to accomplish this: (1) *Datei > Dateikonvertierung*, (2) [Ctrl]+[e] and (3) click on the rightmost icon *Konvertieren*. This opens the conversion dialog. As format, use TIFF (16 Bit) without any other options enabled except the one to select the folder into which the converted TIFF is to be written.

The TIFF images are not linear but have a gamma applied to them. I do not know what exponent is used, it appears that it is about 2.2 but there is no way to change it.

Below are the exposure series taken at f/7.1 and 100 ISO at 1/4000, 1/1000, 1/250, 1/60 and 1/15 seconds (from left to right) and used to calculate the camera curve on the next page.





The cross marks the cross-over point, where the pixel values are not subjected to gamma.

## Sharpening

RAW images are transferred soft from the camera and are intended to be sharpened a bit. The lens may also be a bit soft so sharpening the developed pictures makes sense. Sharpening can be done in Nikon View NX 2 under *Anpassungen > Scharfzeichnung* in the sliding window at right where the colour aberration option is. There is only a 1 to 10 slider and it is difficult to get the right amount but a moderate 3 seems appropriate.

Photoshop offers Smart Sharpen, which gives more control over the function and shows the result immediately. *Filter > Smart Sharpen* opens a dialog.

In this dialog select *Lens Blur*, set *Amount* to 50% and *Radius* to 1.5 pixels. This sharpens the image a bit without introducing halos at the high contrast transitions.

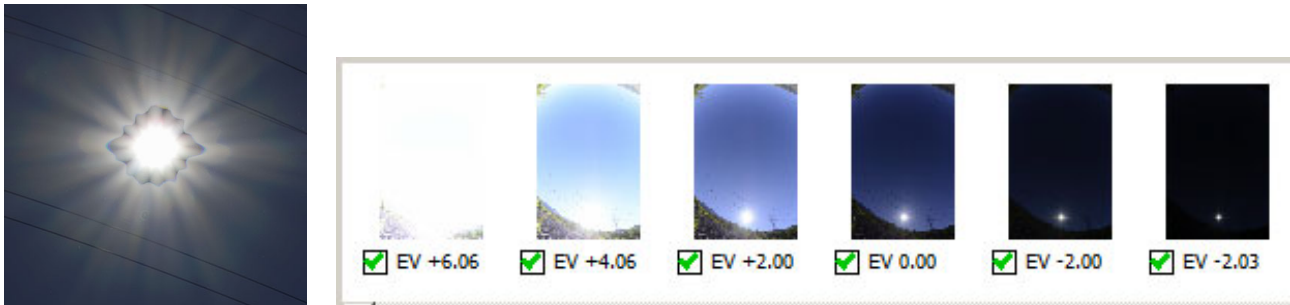
After many tests, I decided to use the sharpen option of Nikon View NX 2 at 3. Subjectively, the result is a bit better than from Photoshop. Apply this filter to all TIFFs.

## Merging LDRI to HDRI

There are several merging programs giving different results. After all the tests performed, I settled to using Pictura 3.2 because it is fast, reliable and easy to use. If the individual shots do not overlap 100%, the result will appear to be less sharp.

## Photoshop CS5

The Photoshop Bridge permits merging an LDRI bracket series to an HDRI painlessly. The process is fully automatic and the individual images are aligned. Select all LDRIs that make up an HDRI, click on *Werkzeuge > Photoshop > HDR pro zusammenfügen*. This opens Photoshop, loads the LDRIs selected, merges them to an HDRI and presents the result. Optionally, *Remove ghosts* can be enabled. It appears that this is the easiest way to merge the files. Unfortunately, there are catches.



Although the exposure data is in the TIFFs, they are not interpreted correctly. The exposition times were 1/15, 1/60, 1/250, 1/1000 and twice 1/4000 second, all at ISO 100. The first five at left f/8 and the last — second 1/4000 — at f/16, this is 2 EV off f/8. The cropped sun above at left shows the result of such a merge. You would not see this after tone-mapping but for a light probe, this is useless.



De-ghosting works surprisingly well for moving twigs but failed to remove a car. After closer examination, the colour of the greenery got greyish and the glossiness got lost. For a panorama, this is not acceptable and living with multiple twigs seems to be the lesser evil

< Left ghosted, right de-ghosted. >



The workflow is absolutely great but it is difficult to get satisfied with the result if you aim for spherical panoramas for computer graphic (CG) light probes.

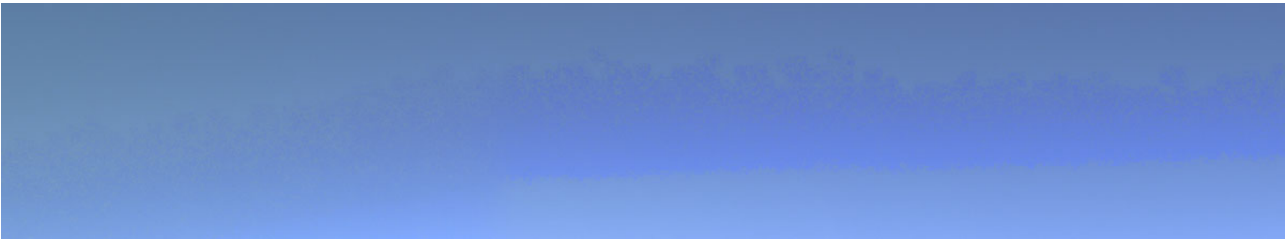
## Picturenaut 3.0/3.2

The workflow with this free program is only slightly more laborious but yields a much better result. *File > Generate HDRI* or [Ctrl]+[g] opens a dialog to load the exposure sequence. The dialogue shows *Aperture*, *Time* and *Bias* correctly in a table if this information can be found in the images; otherwise enter the appropriate data manually.

There are four options that can be selected. *Exposure correction* should be enabled because it does a great job correctly calculating the true EV range. *Automatic image alignment* can prove to be helpful if it does not get fooled by moving targets. Disable it only if you get a funny result. Remember that you get the sharpest result if all exposures are fully aligned.

*Ghost removal* is not efficient and it can mess up the colours. I had a cloudless sky with bands in it (see below) which disappeared once *Ghost removal* was disabled. *Color balancing* should not be used if this is going to be a part of a panorama. The shots from all positions must be treated exactly the same.

Then, there is *Weighting*. *Standard* gives the best result.



The *Curve* needs some attention. I used to leave it at the default *Compute* for the shooting position with the highest dynamic range. Picturenaut 3.0 saves the calculated curve along with the merged HDRI (Picturenaut 3.2 does not so anymore). For the subsequent merges, I used *User defined* and loaded that calculated curve. I was quite happy with the results.

When I acquired the camera curve (see page 8 above) I noticed that the curve is different if colours are involved, which is usually the case. Only a bracket sequence of grey results in a true representation of the curve. I used the sequence above with which I created the curve graph in HDRShop and use it in Picturenaut 3.2 when merging the LDRIs to an HDRI.

I merged bracket sequences to HDRI using the *Standard curve* (gamma 2.2) and my calculated curve. Using the *Standard* curve resulted in a 10% higher dynamic range and warm, slightly yellowish result. Using the curve I calculated resulted in a slightly cooler HDRI with a bit more blue, making the sky more beautiful at the expense of the green on the vegetation. However, these are white balancing issues and can be corrected in the final HDRI panorama.

After much testing, I settled to use the camera curve for each merge. It is really very difficult to tell whether the actual camera curve or the standard curve is better.

### Other Merging programs

I tested some other merging programs but found none is easy enough to be used on a regular basis.

**Luminance HDR** has nice options but it takes ages to merge six 16-bit TIFFs 24 Mpx to an HDRI. The quality is good, though. It complains about wrong TIFF tags several times for each individual file. The TIFFs were saved from Photoshop and even the AsTiffTagViewer shows the contents of the tags correctly. The main issue here is how long it takes to get an HDRI.

**FDR Tools Basic** is a bit awkward to use with all these floating windows but it merges the LDRIs to HDRI quite fast and the result looks good. It takes the exposure data from the TIFF and interprets them correctly. A project has to be created first, and then the exposition series can be loaded. The nice thing is that there is a histogram for each LDRI loaded and one for the merged HDRI. Each LDRI histogram can be adjusted and the result is shown in the HDRI histogram. Preview is also very nicely done. Of course, for a panorama, the histograms must not be tempered with.

**Photoroom (Artizen)** lets you chose the colour space and since I have the camera set to the wider Adobe RGB and not sRGB, this is great. There are four alignment options including a manual one. Then, it sports Merging options like Luminosity Precision for 3D IBL and de-ghosting in an automatic and a manual mode. However, enabling the Luminosity Precision for 3D IBL option gives unexpected results: the blue sky with the sun got white, the sun black. Without this option, the merged HDRI looks quite nice but the centre of the sun is eclipsed: a black disc. The program has an affinity to crash.

**HDRShop 1.03** cannot interpret the camera settings embedded in the TIFF tags and actually complains a lot about unknown tags though the images are loaded. Gamma of the LDRIs can be set from 0.00 to 4.00 and I assume from the camera curve that it should be set to 2.2. The used f-stops have to be set; there is a button for 2 F-stop. Also, the lowest and highest pixels

can be set and only the values in between will be taken into account in each LDRI when merging to an HDRI. However, the range can also be calculated but for a panorama, that would be a bad idea. The merged HDRI is correct and is available quite soon but “features” some very bright single pixel artefacts (fireflies). Earlier tests also showed that very bad colour errors can appear.

### Checking the generated HDRIs

I highly recommend taking the time to examine each individual HDRI generated closely, particularly for ghosts.

### Ghost removal

If there are ghosts on the vegetation because of an evil breeze, you might have to live with it. However, if there are objects, they need to be removed.



The above exposure sequence of 1/15, 1/60, 1/250, 1/1000, 1/4000 at f/8 and 1/4000 at f/16 shows a car only in the exposure f/8 1/4000 exposition. It is not bright here but in the merged HDRI it is very obvious as a ghost.

The car only appears in the 1/4000 f/8 exposure because of its speed of (allowed 80 but more likely) 100 to 120 km/h. It is much simpler to remove such ghosts in the original 16-bit image than in the merged HDRI because the errors introduced are blended in the final HDRI.

Below the same part as above cropped from the tone-mapped HDRI. At left, the series from above merged to an HDRI in Photoshop; centre merged with de-ghosting enabled and at right after manual correction (zoom-in, the images are embedded large enough).



In the case above, I copied a part of the road and its near and far border and put that part over the car — erasing a part of the vegetation in the process. I selected the vegetation in the original image with the magic wand and copied that into the new image without the car restoring the vegetation approximately. Since the picture is very dark, it was quite fiddly to get it approximately right.

There are many strategies to remove unwanted objects and restore the desired ones. Which method is best has to be figured out from case to case. The LDRI series was merged again to an HDRI using the manually de-ghosted image. Remember that in Photoshop the Magic Wand does not work for 32-bit images, only for 16 and 8-bit.

## **Lens flares removal**

These flaws can be easy or nearly impossible to remove. A better lens creates less flares, a neutral density gel filter at the back of a lens (if possible) also works wonders. The Nikon 10.5 mm DX fisheye lens quite generously creates flares, the 0.6 ND Watten gel filter I put in the back was well worth the money and the trouble to cut it to the fitting size. The Nikon 18-15 mm and 16 mm FX fisheye lenses create almost no flares. There is no option to furnish it with a ND gel filter but since it is so much better than the DX lens with the filter, there is actually no need. However, a bit of lens flare is still visible and best removed in the merged HDRI.

The lens flare at left can be easily removed if there are no clouds. Select a part of the sky near the flare and give it a soft edge, 16 pixels or more, copy it and then paste it over the flare.

The rainbow flare on the stone is very difficult to properly remove and it may appear less offensive if left there then if tempered with.

The haze is introduced by the sun lighting up the humidity in the air and the lens is not responsible for it. Enhancing contrast can dim the effect but that is out of question for an HDRI panorama for CG use.



## **Panorama Stitching**

Panoramas can be stitched using directly the HDRI but this is not the most efficient way. Rather, tone-map the HDRI series to LDRI and use those. Even simpler and giving quite good results are selecting the hero exposition of the brackets and stitch them. Once happy with the panorama just replace the LDRI with the HDRI and stitch again.

## **Hugin**

Hugin is a great free panorama stitching program. I used it in the beginning but got a bit lazy setting control points, I moved on to PTGui, which looks similar to Hugin, is mostly the same under the hood but automates the control point generation and stitching process nicely.

## **PTGui Pro**

This is the tool to stitch spherical panoramas as painless as it can get. To do HDRI stitching, the Pro version is mandatory. To load directly the HDRI segments to stitch the panorama is not the fastest and most reliable way as I found out. I did this in the beginning on a computer with only 2 GB of RAM using 6 Mpx HDRIs. It took very long to complete the panorama and while waiting, I discovered more than 1000 temporary files occupying 30 GB on the hard disk. Obviously, with so scarce memory, everything has to be bled onto the HD and gathered back from it. This is what makes the process so long.

The alternative is to use LDRI, either tone-mapped HDRI or using the best exposed LDRI of the bracket series. With the latter method, stitching may not yield the best result and another

attempt must be made with the next faster exposed series. Using the hero exposed images is faster if it works; using tone-mapped HDRI is slower but works more reliably.

The best results to stitch tone-mapped HDRIs I usually get when tone-mapping with the Reinhard tone-mapping operator plug-in in HDRShop 1.0.3 (works in Picturenaut, too). It is fast and gives consistently good results. I wrote myself a small tool to tone-map a whole batch of HDRIs automatically. For night shots, tone-mapping is not recommended, however.

**Lens Database:** I had created a part panorama with different lenses, stitched and optimised them and saved the lens settings in the data base. This data base is a pure text file and can be edited with any text editor. It is a bit difficult to find. In Windows 7 look under:

```
C:\Users\<yourname>\AppData\Roaming\PTGui\lensdb.ptl.
```

Editing this database makes particularly sense to sort the lenses stored in it. In PTGui, the data base can be accessed via *Advanced > Lens Settings > Lens database*. PTGui can interpret the lens data from the TIFF/EXIF — provided the 16-bit LDRIs are used to stitch the panorama — but loading the lens data saved in the database gives a much better preliminary result. The panorama is almost done. If the tone-mapped HDR images are used, there is no lens data contained and the parameters have to be entered manually anyway. So using the lens database is in any case the best option. Deselect *Automatic* for the *Camera/Lens* parameters.

**Load the images:** From the *Project Assistant* tab click on *Load images* and select the series that will make up the panorama. The images are loaded as thumbnails and can be rotated with the controls far right of the thumbnails. Good practice is to rotate the RAW NEFs already in Nikon View NX 2 before converting them to 16-bit TIFF then you do not have to worry about that anymore. Load also the two nadir shots, but not the one with the camera offset.

**Check loaded images:** Select the *Source Images* tab and check if all are loaded in the correct order. There are buttons on the lower left to put them into a different sequence. However, if the pictures have not been renamed, the number given by the camera sorts the images already in the right order.

**Access Advanced Options:** From the *Project Assistant* tab, the *Advanced* button can be found at far right. This opens a plethora of additional tabs.

**Set the Panorama Type:** Open *Panorama Settings* and select the projection from the drop-down. This will usually be *Equirectangular* for a spherical panorama. Then set the horizontal field of view to 360° and the vertical one to 180°.

**Set Image Parameters:** Open the *Image Parameters* tab and enter the rotation for each picture. For a fisheye sequence, this will be for the horizontal shots Yaw 0, 60, 120, 180, 240, and 300 degrees. *Pitch* and *Roll* remain at 0. For the zenith shots, it would be Yaw 0 and 90, *Pitch* 90, and -90 for the nadir. If there is more than one nadir, adjust Yaw accordingly. *Roll* should be left at 0 for all images. Right of each image thumbnail is a thumbnail of the full panorama and the images shown at the appropriate location within the panorama. If the zenith appears at the bottom, *Pitch* was set wrong.

**Select the Lens:** Open the *Lens Settings* tab and click on *Lens database*, then select the lens used to acquire the panorama. If the *Panorama Editor* is open, you will see how everything falls in place.

**Align Images:** Back at the *Project Assistant* tab, click on *align images* and wait until the control points are generated. You might now look at the control points listed in the table with the highest error. Just double click on a value to open the *Control Points* tab. Such control points are mostly on twigs, foliage or shadows that moved or on clouds, if there are any. Just delete them. At least 4 control points should be left.

**Define straight lines:** In each image, define some vertical and horizontal lines if there are any. In landscape shots this may prove to be impossible but in urban environments or inside, there are many. Select the same image for the left and right window to add the lines. Such lines can greatly improve the accuracy of the stitch.

**Run Optimiser:** In my experience, running the optimiser first for only the horizontal shots — the cylinder — is a good strategy. Then add the zenith and nadir. Use the *Advanced* options. Make sure *Field of View*, *a*, *b* and *c* is ticked for *Yaw*, *Pitch* and *Roll*, except for the reference image. Optionally enable *Vertical* and *Horizontal* shift. Run Optimizer.

Add the zenith and nadir images in *Viewport* and run the optimiser again. In my experience, there is nothing wrong to have these images in the *Yaw/Pitch/Roll* columns and not in the *Viewport*.

Check the *Control Point Table* for badly matched control points. If there are, check where they are and either delete or adjust them. Run the optimiser again with another source image. This can improve the result. After three or four iterations, the stitch will be announced as very good. Have a look at the image parameters. The higher the values in the *Roll* columns, the sloppier the tripod had been levelled. Save the project at this point.

**Add the offset nadir:** Only if you made an additional nadir shot to photograph the space initially covered by the tripod. Set the camera orientation (*Yaw* and *Pitch*) as appropriate and connect it with manually set control points to the nadir shots. Keep the control points on the ground; do not try to match parts above the ground.

**Mask the Tripod:** Open the *Mask* tab and select one of the two nadir shots. Use the red marker to paint out the offending parts; then select the second nadir shot and do the same.

**Continue with offset nadir:** In *Optimizer* (*Advanced*), enable *Viewport* for this last image and optimise. This usually works quite fine, if not, check

<https://www.ptgui.com/examples/vptutorial.html>

and follow the procedures towards the end of the page. In the *Image Parameters* tab, towards the right edge, there is a column *Blend priority*. Reduce it from 100% to any fitting value down to 1%, 25% is a good start. This last shot is usually the worst to correct and stitch.

**Optional:** At this point, you may also want to shift the panorama to get the most important part into the centre. In a landscape, that would be the sun. Straighten the panorama and optimise again if the panorama was shifted.

**Create LDRI Panorama:** Click on the *Create Panorama* and set optimum size to *Maximum*. The «Ken Turkowsky rule» states that only 70% of the maximum size of a panorama can be used. This is certainly true for large prints. Using the final panorama as light probe and optionally as (tone-mapped) backdrop in CG this can be looked at a bit more relaxed.

Check the full size panorama for stitching flaws once it is created. If there are any, attempt to correct them by introducing additional control points, optimise and create a new panorama.

**Save Project:** If everything looks good, save the project from *File > Save Project As* or with the shortcut [Ctrl]+[Shift]+[s]. The file created has the .PTS extension and it is a pure text file that can be edited in a text editor if so desired.

**Swap the Source Images:** In the *Source Images* tab, each individual LDRI can be replaced by the corresponding HDRI. Alternately, the project.PTS file can be edited by replacing the file names. For a panorama shot with a rectangular lens that contains many images, this might prove to be a handy option, just search/replace the file names. However, if a fisheye lens was used, replacing 8 to 11 images is swiftly done in the GUI.

**HDRI:** Open *Exposure/HDRI* tab and *Enable HDR* stitching and select the Method *True HDRI*. Though I do not enable *Exposure* correction, I have the Vignetting and Camera response curves optimised and I also save the curves. I do not fine tune exposure and white balance. These things are done with the final HDRI panorama.

Open the *Create Panorama* tab and select *HDRI*. PTGui has a nice tone-mapper but I usually do not use it and disable the LDRI output, only generating the HDRI panorama.

**Check HDRI Panorama:** When the panorama is stitched, which takes a moment for a 16k HDRI panorama, open the result in HDRShop and check exposure steps and stitching flaws. If you are happy with the result then save the project with a new name and close PTGui.

## Inserting the Nadir and removing the Tripod and Shadows

This can be quite a tricky and time consuming task. There are also several methods how it can be done; one for the nadir is using PTGui as mentioned above. The simplest method is to put a mirror ball on the ground to cover up the missing nadir. At rare instances, this may also be the last resort to make a spherical panorama whole.

## Photoshop CS5 and Flaming Pear Flexify 2

This is the most cumbersome and time consuming method I have tried so far. Open the HDRI panorama and copy the image by pulling the Background Layer over the New Layer icon at the bottom of the layer window (second last at right). With this layer selected, start the filter Flexify 2. Select Equirectangular for input (first in list) and Zenith & Nadir (way down in the unsorted list) for output. Keep all values at 0. Select Transparent gaps and De-halo, click on OK and wait. Open to Task Manager to confirm that Flexify is doing something. Eventually, Photoshop returns and loads the filtered image and shows a progress bar.

Once the nadir is inserted — we will come to that a bit farther down — open the Flexify 2 filter again, then swap input and output format. The list for input projections is not so long like the one for the output projections, so Zenith & Nadir is swiftly found and Equirectangular is on top of the list. Nevertheless, it is a good idea to save those transformations in a memory dot.

It takes again ages until Flexify closes after OK is clicked and Photoshop reappears, reading the transformed image. This is again a spherical panorama but contains only zenith and nadir. Select the nadir in the image, [Ctrl]+[c] it and [Ctrl]+[v] it into the Background layer. The copy can be deleted and the corrected spherical panorama saved.

## Pano2VR 64-bit (version 4.5)

The main aim is to create virtual reality movies, either as QTVR, Flash or HTML5/Javascript. But Pano2VR can do more. It can extract the cube faces from a spherical panorama or assemble a spherical panorama from individual cube faces and related projections like crosses, strips, cylinders. And it does LDRI and HDRI.

**Patch:** It has yet another trick up its sleeve: it can patch the panorama. This is great if only one part in the panorama needs attention, i.e. the nadir. Load the spherical HDRI panorama, click on Patch Input to open the Patch Panorama dialogue. Click on Add, then pull and drag the panorama in the window until the patch to be corrected is in the centre. You can also enter the Pan (Yaw), Tilt (Pitch) and Roll numerically, even the field of View (FoV).

Once the image shows where it needs attention Extract it and save it in the same format as the input panorama is. Then correct this part externally and save the corrected image with a new name. In the Patch Panorama dialog, Load the corrected image, Insert it and click OK. The patched image is inserted in the panorama, replacing what was there before.

This is a very cool option and really fast. This method permits accessing any part of the panorama interactively to correct flaws. The problem may be rather that the preview cannot show small issues. To really scrutinise the panorama and correct even small parts, another method can be used.

**Extract all cube faces:** Click on Convert Input and select Cube Faces as output, make sure to keep the file type the same, e.g. hdr. You may also want to increase the size of the cube faces by one or two pixels. This converts the spherical panorama to a cube and saves the six cube faces. They can then be opened with an external program and corrected. Once done, just Select Input and go for the Cube Faces. Enter the file name for each face and once loaded, Convert Input again selecting Equirectangular this time. The size of the panorama can be set to match the original one. The only nuisance is that you have to navigate all the way down the folder hierarchy to access the corrected cube face images and do so for each one. The cube faces are identified as 0 = north, 1 = east, 2 = south, 3 = west, 4 = zenith and 5 = nadir. Best save the corrected faces with retaining the number in the filename somewhere. I usually add a «c» at the end of the name so that myname\_cube\_0 gets myname\_cube\_0c. This makes identifying the files easier and can be typed directly in the open dialogue. Version 5 works a bit differently.

### Actually correcting the nadir

Whichever method was used to extract a part of the panorama to an undistorted image, that image has to be processed. Photoshop seems to do the job best even though a lot of effort went into making this a bulky piece of software that also changes the shortcuts from version to version. Giving a selection soft edges works with [Ctrl]+[Alt]+[d] in CS3 and with [Shift]+[F6] in CS5.

The Magic Wand does not work for 32-bit images and this is really annoying. It would have been so easy to just select the black nadir with the magic wand. Instead, the Lasso has to be used, which is tedious to select the tripod that was masked out in PTGui. If no nadir shots were made, a circle can be drawn.

Now comes the cool part. Right click into the selection drawn and from the context menu, chose Fill. A new dialog appears where you can set Context Aware (only in CS5, this is missing in CS3). Use Normal for Blending Mode and Opacity 100%; do not enable Preserve Transparency. It will take a moment and the black nadir gets populated. If the selection gets near other objects than just the floor, this will be also considered as viable content and it messes things up a bit. Usually, the result is quite acceptable, though. You can always select a smaller part and Fill that one the same way to correct any misplaced parts.

A more tedious way is to get parts from around the nadir hole and paste them into it. It works fine on a paved ground or on a grassy or scree patch but gets very difficult for a tiled or wood plank floor.

If there is an offset nadir bracket sequence, each LDRI can be corrected with PTLens, either the Photoshop plug-in or the stand alone program. PTLens corrects the parallax errors but unfortunately, is not HDRI compatible. Setting each shot of the bracket series exactly the same is difficult. Actually, PTGui does the same as PTLens when *Viewport* is enabled, and can do HDRI. In fact, I discovered this when I was fiddling with PTLens.

### Inserting a Mirror Ball to fill the nadir

If there is no nadir shot or the nadir cheated in looks too cheap and obvious, it may be better to resort to yet a cheaper method but one that looks at least clean. Inserting a mirror ball is easy and straightforward.

Load the spherical HDRI panorama into any 32-bit capable graphics application, e.g. Photoshop. From the lower edge upwards, select a rectangle that just covers everything black. Adjust the selection with pixel precision. [Ctrl]+[c] it and create a new file to get the vertical size of the selection but do not actually create it. Invert the selection to get the upper part of the panorama.

Copy this one and create a new image and paste the panorama in it. Then resize that image but without having height and width linked. Just squash it vertically to the size of the lower black part. Then flip this image strip vertically (Image > Image Rotation > Flip Canvas vertical). Copy this image [Ctrl]+[c] and paste [Ctrl]+[v] it into the spherical panorama with the missing nadir, move it to the lower edge and flatten the image ([Ctrl]+[e]). That is all there is to it.

## **The Shadows**

We still have not covered the shadows. They are best eliminated from the cube faces because there are no distortions. Extract all cube faces with Pano2VR. Open each cube face that has some undesired shadow from the tripod (or yourself) in it. The best strategy here is also to draw a line around it using the lasso tool and fill the selection context aware. Also copying a part left or right of the shadow with a soft feathered edge and put it over the shadow works like a charm. Sometimes, the clone stamp is good enough to do the job. For an outdoor, the lower the sun stood, the longer the shadows.

Do not forget to hide your steps in the snow when you do a winter shot, or in the sand.

## **Other issues to correct**

To eliminate dirt on the Bayer filter that clutters the sky, lens flares and such nuisances, copy part of adjacent content with a soft feathered edge over the speck or blotch. The clone stamp can also be of great help for small spots.

Moving clouds are difficult to correct. The edges may look as if the cloud was made of a number of misaligned layered sheets. If there are only few of those buggers, handle them like the tripod shadows. Using a Gaussian filter also helps but if the selection edges are too soft, a blurred part from the environment may bleed into the sky, making things look much worse. The wind may cool you down when shooting the panorama in the sun but it may heat you up when you have to remove moving clouds.

If the panorama is to be used as light source for image based light (IBL), putting the main light source, the sun for example, in the horizontal centre is a good idea. This makes all your light probes consistent. If it is used as a poster, the main focus needs to be moved to the correct location. In either case, this can be done in HDRShop: Image > Transform > Shift w/wrap and enter the number of pixels to shift it right, what gets pushed out at the right comes back on the left. Another free program that can rotate HDRI and LDRI panoramas is PanoRot, a free application I wrote.

## **White Balancing**

Shooting panoramas with the camera white balance set to Full Sunlight is usually the best strategy. If you are inside, fluorescent and tungsten lights may be mixed making it difficult to get the right choice for the white balance setting in the camera. Keep it to the sun since you will have to adjust white balance anyway in many cases.

If a neutral white or grey paper is visible somewhere in the panorama that is your ticket. If not, look for white walls, grey stone, snow on a mountain or whatever neutral coloured object you find in the panorama.

I am using HDRShop for white balancing. For large panoramas, the program has to be made large address aware (LAA) in order to be capable of loading it without crashing. Select a large or small rectangle with a colour that should be neutral. Then, File > Pixels > White Balance Selection does the trick quite painlessly and fast.

In Photoshop, you can experiment with Image > Auto Color. It takes quite a while but the result is usually surprisingly good.

## Light Probe Generation for IBL

Most programs can use the HDRI in the spherical projection, Bryce is among them. However, I used to prefer an angular map for Bryce and transform the spherical projection to the angular map one using HDRShop but not anymore. In order to get fine lines properly, at least 3 x 3 super-sampling is needed and this makes the result soft — or a tiny bit blurred.

Bryce is one of the 3D programs that do not use the Monte Carlo algorithm to get the light and colour information randomly from the light probe, but the median-cut one. The practical upshot is that it can use a high resolution HDRI for the light without creating the fireflies and noise known from the Monte Carlo method. The disadvantage is that it creates banded shadows, though these can be softened — at a heavy render time penalty. The Monte Carlo algorithm prefers a specular or even diffuse convolved HDRI for the light. This means that it cannot be used for the reflections, for which a second HDRI is needed and then, yet another LDRI panorama in high resolution for the backdrop.

Bryce can use a high resolution HDRI for the light, for the reflection and can even tone-map it for the backdrop and if a blurred HDRI is needed, can diffuse and specular convolve the HDRI up to a Phong exponent of 100. On the other hand, Bryce being still a 32-bit application does not have a lot of memory left for a truly high resolution HDRI. However, Bryce can be set large address aware (LAA) and can then handle up to 3.5 GB instead of only 2 GB.

## Cleaning up

Create the tone-mapped LDRIs for internal and public publishing and for your Light Probe database. From the tone-mapped spherical panoramas create QTVR, Flash and HTML5 in different sizes for publication. Pano2VR can create such movies out of the HDRI panorama. Tone-mapping is not the best as it could be but may be acceptable in some cases.

Delete unwanted originals, the TIFFs and all intermediate image files. It is a good idea to record what you have done in a simple text file. Also identify the images, numbers are not very helpful. Here is an example how this could look like (16 mm Fisheye):

	-000°-	-060°-	-120°-	-180°-	-240°-	-300°-	-Z 0°-
0.066667s	1519	1526	1533	1541	1549	1558	1512
0.016667s	1520	1527	1534	1542	1550	1559	1513
0.004000s	1521	1528	1535	1543	1551	1560	1514
0.001000s	1522	1529	1536	1544	1552	1561	1515
0.000250s	1523	1530	1537	1545	1553	1562	1516
0.000250s	1524	1531	1538	1546	1554	1563	1517 f/16

Pack this all together and save it to an external drive. There are images you will never be able to reshoot at a later date so it is good practise to save the originals. From them, you can always make the panorama anew if you are more experienced.

## **Panorama Downscaling**

The final HDRI panoramas are usually too big for image based light application in a 3D program because of their memory footprint and need to be downscaled. The most popular HDRI file format is the Radiance HDR. The original image data was 96 bit: 32 bit as single precision reals for each colour red, green and blue. The HDR file format is actually only 32 bit: 8 bit each for the mantissa values for the three colours and one byte for the common exponent which is taken from the brightest colour. The colour plains and exponent are lossless compressed by a run length encoding method. The file size does not show how much memory the HDRI actually uses when loaded. The memory usage is rather width x height x 12 bytes.

## **Limitations for HDRI**

If an LDRI must be downscaled, a good strategy is to first downscale it to double of the final size. Then apply a sharpening filter and downscale to the required size. There are several methods to sharpen; doing it in the frequency domain applying a FFT (fast Fourier transform) filter usually gives nice results. Transitions between brightness and colour are found and made narrower.

It appears that there are no such filters for HDRI. Convolving high-pass filters, unsharp masking and their derivate — like smart sharpening — work on the image plane, not its frequency. The contrast between transitions is enhanced and thus gives the impression of more sharpness. This works quite fine but on already high contrast transitions, black artefacts are introduced.

On lines that are not straight, not vertical or horizontal, anti-aliasing gets lost and the lines become staircases. This is true for all sharpening filters, the one found for HDR images make it worse because the steps are yet enhanced by artefacts.

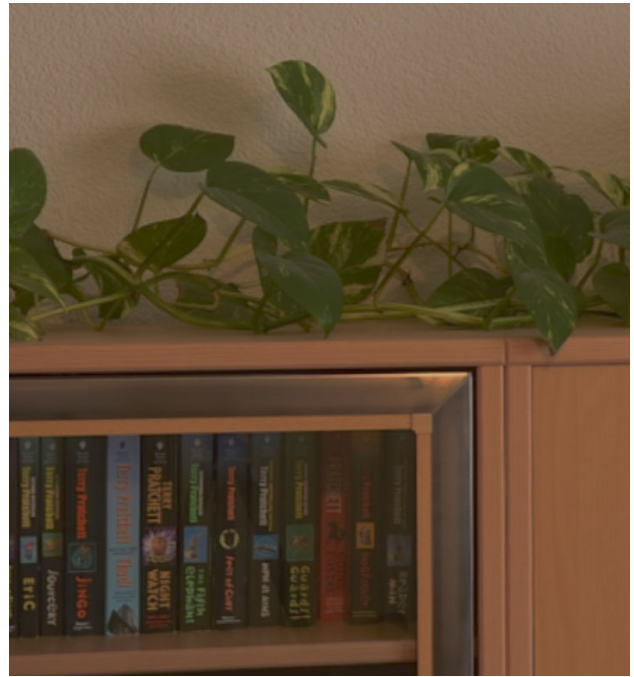
Because of the filters available, sharpening HDR images before downscaling is not a good idea. Downscaling is usually done by resampling the image using bilinear interpolation. To enhance anti-aliasing, super-sampling should be used when available. However, super-sampling blurs the image; the higher it is set, the more blurred the result gets. Super-sampling 1 x 1 is like none, the next step is 2 x 2. This appears to be an acceptable compromise when a spherical HDRI panorama is downscaled; for angular maps use 3 x 3 super-sampling.

Below are examples that show the aliasing and artefact issues when sharpening an HDRI and also the loss of sharpness when super-sampling is introduced when downscaling or remapping.

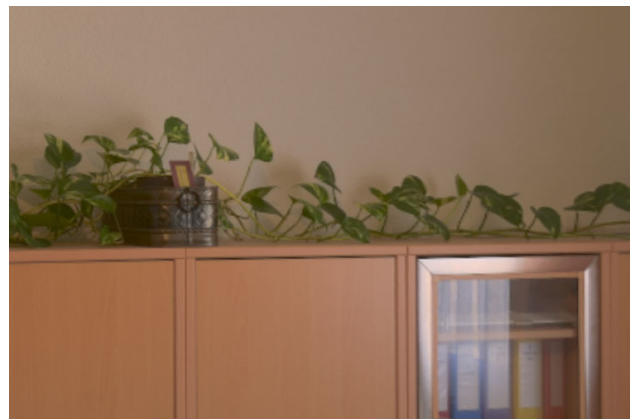
## **Examples**

All examples are Bryce 7.1 renders and use the HDRI tone-mapped by the Photographic operator by Erik Reinhard as implemented in Bryce 7.1 pro.

## Creating HDRI Panoramas

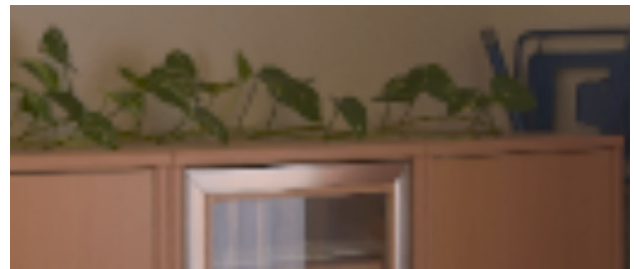
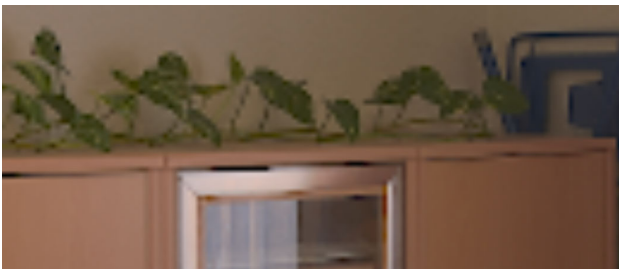


Above at left, smart sharpened with a radius of 1.5 px and 50% before downscaling, at right without sharpening and with 2 x 2 super-sampling. The left image is sharper and the introduced artefacts almost invisible.



The introduced artefacts and the poor anti-aliasing is obvious in the left render, even though the right one is a tad less sharp, it looks better in my opinion.

Below, HDRI half size, this becomes very obvious at the skewed lines.



The examples below show how bothersome and obvious artefacts can get at high level transitions. This is simply not acceptable anymore. In fact, I got aware of the problem when I was about using this HDRI as light probe for an artwork.



The artefacts in a slightly sharper backdrop are more annoying than a background that is a bit on the soft side.

The more the original HDRI is downsampled, the coarser the artefacts get. The first two rows of examples shown were downsampled to 8192 x 4096, the second two to 4096 x 2048. The example with the window above is repeated below but downsampled to 2048 x 1024. There are not only dark artefacts on the window, they can also be found on the loudspeaker and the lens on the table.



The dark artefacts also appear when the HDRI is not tone-mapped and rendered with and without gamma engaged.

### How to downscale

When converting the raw 14-bit NEFs to 16/48-bit TIFFs, sharpening 3 (of 10 in Nikon View NX 2) is good, then merging them to HDRI (with Picturenaut 3.2 using the D600 camera curve) and create the panorama (in PTGUI). The raw panorama can be converted into the six sides of a cube (in Pano2VR) then corrected and retouched (in Photoshop), transformed back into the spherical projection (Pano2VR) and the moderately sharpened (smart, 50%, 1.5 px in Photoshop). However, if the HDRI is used at the full size, do not sharpen.

The pixel values can then be adjusted (in HDRShop) so that the output level of the light probe gets standard in Bryce 7.1 Pro, the panorama shifted to have the key light at azimuth 0°.

Do not use *Image > Scale* in HDRShop, the result has no artefacts but anti-aliasing gets very bad. Use *Panoramic Transformations* spherical to spherical, bilinear and 2 x 2 super-sampling. Use the original HDRI for all downscaling. Do not downscale a downsampled HDRI.

Photoshop can downscale by bilinear resampling; the result is the same as in HDRShop if 1 x 1 super-sampling is used.



Here, the original 16,112 x 8056 px tone-mapped HDRI is shown without downscaling showing good anti-aliasing and no artefacts.



Downscaling by resampling using 2 x 2 super-sampling the original spherical HDRI seems to me the best compromise to start with. Some sharpening may be possible, depending on the HDRI and the scene depicted.

## Links

Here is a list of links in alphabetical order pointing to stuff mentioned in the text above. Note that the free HDRShop 1.0.3 is not available anymore and for Fhotoroom a login is needed.

## Software and hardware mentioned

Bryce: <https://www.daz3d.com/bryce-7-pro>

FDR Tools: [http://www.fdrtools.com/front\\_e.php](http://www.fdrtools.com/front_e.php)

Fhotoroom: <http://www.fhotoroom.com/products/tutorials.asp>

Flexify: <http://www.flamingpear.com/flexify.html>

HDRShop: <http://www.hdrshop.com/>

Hugin: <http://hugin.sourceforge.net/>

Luminance HDR: <http://qtpfsgui.sourceforge.net/>

Manfrotto: <https://www.manfrotto.us/> <https://www.manfrotto.co.uk/>

Nikon: <http://www.nikon.com/>

Pano2VR: <https://ggnome.com/pano2vr>

Panorama Studio: [https://www.tshsoft.de/en/panostudio\\_index](https://www.tshsoft.de/en/panostudio_index)

Photoshop: [www.adobe.com/products/photoshop.html](http://www.adobe.com/products/photoshop.html)

Picturenaut: <http://www.hdrlabs.com/picturenaut/>

PTGui: <http://www.ptgui.com/>

PTLens: <http://www.epaperpress.com/ptlens/>

## Vanity Links

Additional reading and images concerning the making of HDRI and panorama on my web site: Go to Raytracing > Panorama? | Mirrorball? | IBL/HDRI? | HDRShop?

<https://horo.ch/>

Some of my tone-mapped HDRI panoramas, Flash and HTML5 as well as some tips:

[https://horo.ch/photos/intro\\_en.html](https://horo.ch/photos/intro_en.html)

Free light probes in the angular map projection (also some made with a mirror ball).  
Go to Resources > HDRI and HR-HDRI:

[https://horo.ch/raytracing/intro\\_en.html](https://horo.ch/raytracing/intro_en.html)

Commercial light probes:

<https://www.daz3d.com/horo/>

PanoRot: <https://horo.ch/docs/progs/zip/PanoRot.zip>

(3.3 MB incl. Doc. Windows 2000 und newer).

The program PanoRot.exe version 1.04 (88 KB) mirrors and rotates spherical HDRI and LDRI panoramas in 1° steps left and right. HDR (RGBE rle-2) and 96-bit TIF HDRI as well as 48-bit and 24-bit TIF LDRI in almost any size can be processed (tested up to 302 mega pixel). The illustrated 6 page documentation is included.

## Literature

Here are a few books that can be recommended for further reading.

### German

*DRI und HDRI* by Jürgen Kircher; ISBN 978-3-8266-5903-4 (2008)

*Das HDRI-Handbuch* by Christian Bloch; ISBN 978-3-89864-430-3 (2008)

*HDRI in der Praxis* by Jack Howard; ISBN 978-3-89864-580-5 (2009)

*HDR-Fotografie* by Reinhard Wagner; ISBN 978-3-7723-6470-9 (2010)

### English

*High Dynamic Range Imaging* by Erik Reinhard, Greg Ward, Sumanta Pattaniak und Paul Debevec; ISBN 978-0-12-585263-0 (2006)

*The HDRI Handbook 2.0* by Christian Bloch; ISBN 978-1-937538-16-3 (2012)

*Digital Lighting & Rendering* by Jeremy Birn; ISBN 978-0-321-92898-6 (2014)

October 2013/horo

**Updated** December 2015, March 2016, January and September 2017/horo