

## Using V-Log on the AU-EVA1

(excerpted from “A Guide To The Panasonic AU-EVA1 Camera”)

The AU-EVA1 includes a logarithmic gamma mode, called V-Log. It is, in fact, the same V-Log gamma as employed in Panasonic’s premium VariCam 35. Using V-Log gives the user the most power over the image for post-processing and grading, when recording internally or to a video recorder.

Many cinema cameras offer logarithmic gammas. A logarithmic gamma is a different way of storing the raw sensor’s brightness information; it isn’t magic, but it is a more post-production-friendly storage method that provides for preservation of the maximum latitude and dynamic range, while also providing more flexibility in adjusting brightness levels in post while minimizing the prospect of “banding” in the shadow areas (the darker areas of your images). V-Log is not just a gamma, it is a “mode” — when you enable V-Log (by choosing SYSTEM SETTINGS>COLOR SETTINGS>MAIN>V-LOG), the camera will disable the entire SCENE FILE SETTINGS menu. Nearly all image processing is disabled in V-Log mode; the camera bypasses all that internal processing, and delivers a smooth, flat, clean image that is ready to be graded in post. You can, however, still assign noise reduction using the camera’s NR menu.

Logarithmic gammas provide for more power in post, but as a famous webbed superhero was once advised, “with great power comes great responsibility.” A logarithmic gamma is not the same thing as a raw sensor image, but there are some similarities in the amount of post-production processing that needs to be done between a raw image, and a logarithmic gamma image. Raw camera sensor images undergo a lot of image processing before they can be said to look pleasing. The raw data off of a camera sensor does not look pleasing at all to viewers. It needs to be de-Bayered (or “demosaicked”), most raw footage needs to be gamma-corrected from raw linear sensor data into a monitor-friendly gamma, it will need to be graded for pleasing contrast and tone, it will need noise reduction applied to it, and it will need sharpening, and it will usually need to be converted into a video-compliant format (meeting the EBU and/or ATSC specifications for UHD or HDTV). All that processing needs to happen to the signal, one way or the other. When the camera’s MAIN COLOR is set to one of the SCENE FILE settings, the camera supplies all that necessary processing. When the MAIN COLOR is set to V-Log, the camera does a bare minimum of that processing; the camera will de-Bayer the footage, and convert the linear brightness data into the V-Log logarithmic gamma, and it will store

the footage in an EBU/ATSC compliant format. The images will be recorded or output in a format suitable for grading, but that's where the hard work starts. You can bypass (almost) all the camera's internal processing, but you will have to replace it with post-processing.

In some ways it's much easier to shoot in V-Log than it is to use the camera's Scene File menu system. For one thing, you can bypass learning about what all those various Scene File menu items do, because V-Log disables them all! It really does mean that the job of "painting" the image is not needed on set; all painting of the image gets delayed to post-processing. When shooting in V-Log, pretty much all you need to do is get a proper white balance, and get proper exposure, and then shoot. Just please understand, there's far more to imagemaking than just shooting — if you were to hand an average client V-Log footage without grading, they would probably be very disappointed in it. The footage needs extensive processing before it is ready for viewing, and professional colorists can charge several hundred dollars per hour to work on footage. Using V-Log doesn't mean less work, it just delays when that work needs to be done, and shifts the burden from the camera operator (or Digital Imaging Technician) over to the editor or colorist.

### When is V-Log The Right Choice?

V-Log is probably most suitable for shooting cinema, drama, music videos, and other types of footage where extensive post-processing, color correction, and "stylized" video are the expected final result. When you know that the footage is going to be extensively post-processed, V-Log provides the flattest exposure and the flattest color palette, and the most dynamic range possible. It really is like raw footage in several ways. V-Log is the right choice when you will be controlling the post-production process yourself, and when you understand and accept the responsibility for performing all the necessary post-production steps that will turn V-Log footage into viewable footage. And, of course, V-Log is the ideal choice when you're using your AU-EVA1 in a multi-camera environment where the other cameras are also running V-Log (as an example, when your EVA1 is operating as the B-camera to a VariCam Pure running as the A-camera).

### When is V-Log The Wrong Choice?

I would suggest that V-Log would be the wrong choice for any scenario where heavy and extensive post-production is not expected or is not the norm. For example, I would say that V-Log would probably be the wrong choice when shooting sports, live events, concerts, conventions, news,

any live broadcast, or any scenario where you are expected to turn over the unedited master footage to the client. Unless your client specifically requests V-Log footage, they probably don't want V-Log and may not necessarily know what to do with it. The camera offers extensive image processing and more power and control over the image processing than perhaps any prior Panasonic handheld camera; when you understand and use that power, you can deliver gorgeous footage right out of the camera, and that's the kind of footage that most clients are likely expecting. V-Log does not look good right out of the camera, it looks flat and pale and muted, because it is not designed for viewing directly, and is not the kind of footage that most clients will be expecting. Furthermore, when a job requires a rapid turnaround from shooting to delivery, V-Log may not be the right choice in that scenario. Post-processing takes time. If someone were shooting a wedding and expected to deliver a same-day edit, shooting in V-Log may be problematic in that the post-processing rendering process might slow down your ability to deliver finished footage.

## Post-Production Processing

There are several crucial aspects of image processing that are necessary to turn raw (or V-Log) footage into pleasing, monitor-ready final images. It seems obvious that we need to color-correct, and to gamma-correct (usually by applying an S-curve to make the images really "pop" on the monitor), but that's not all. Many producers have taken those steps only to have people may look at the resulting images and complain that they're "noisy" and "soft." Post-processing is not strictly about color and contrast! Sensor images have always needed noise reduction applied, and the EVA1 is no exception. The EVA1 can do noise reduction internally when it is normally processing footage, and includes several menu options for controlling sharpening and noise, but many of those processes are bypassed when recording V-Log. You can still use the camera's CAMERA SETTINGS>NR noise reduction choices, but you won't have access to the detail level or detail coring settings or skin tone detail features. The camera may still be doing some noise reduction, but without the benefit of user controllability of how it performs that task, the job may not be done to your satisfaction. If you are bypassing the internal processing by using V-Log mode, you should expect to do noise reduction on the recorded V-Log footage or, yes, the footage may indeed look noisy, especially in the darker shadows and lower midtones. Likewise, video footage generally benefits enormously from some manner of sharpening. The EVA1 provides options for sharpening the footage, including the Detail and Coring controls. If you bypass that processing by using V-Log, then you're likely going to need to apply some

sharpening in your post-production phase or, yes, the footage may look soft. The EVA1 is an extremely sharp camera, but video needs some degree of edge enhancement (sharpening) to really make that resolved detail apparent to the viewer. As always, the level of sharpness you add should be governed by the anticipated size of the final display device; the smaller the screen your footage will be viewed on, the more sharpness you can add to the footage in post. The larger the display (such as a movie theater screen), the less sharpening you would want to add in post.

## Exposing For V-Log

V-Log distributes the EVA1's 14 stops of dynamic range logarithmically (or, for the brighter stops, that means it distributes it roughly equally) in its recorded format. Other gamma curves do not, and this means that you have to expose somewhat differently for V-Log. Most video camera gamma curves are designed to replicate what-you-see-is-what-you-get when the footage is displayed on a video monitor. They're convenient, but there are sacrifices made in the amount of dynamic range that can be preserved when using a conventional video gamma. Furthermore, video gammas are more linear in design, and that means that more "bits" are allocated to storing the brightest stops, than are allocated to store the darkest stops. This can result in mushy shadows, and noise and banding in the shadows if you try to brighten them in post. With V-Log, the bits are allocated more evenly, with approximately the same number of shades of gray allocated to each of the midtone and highlight stops; and while the darkest stops do have fewer shades of gray than the midtones might, they generally have significantly more than in a conventional video gamma.

This re-allocation of bits means that V-Log is suitable for recording, but it is not all that suitable for viewing. It simply doesn't look "right" when displayed on a video monitor; it looks very flat and muted (unless, of course, your monitor is capable of applying a Look Up Table (LUT) to the footage.) As mentioned before, V-Log requires post-processing to convert it into something that will look suitable on a computer monitor or television. You can always configure the HDMI or SDI outputs to overlay a V-709 LUT on their monitor outputs, so that people viewing the live feed will see a monitor-friendly version, but that V-709 LUT may or may not represent your final post-production decisions. The V-709 LUT is good for making the video feed client-friendly, but it is not necessarily representative of what the final footage will look like.

Unlike the video gamma curves in the EVA1, V-Log is designed to mimic the characteristics of a film negative that has been scanned digitally. It is

not designed for monitoring and exposure, it's designed to provide a broad, flat scene that can be manipulated in post production. While the camera itself can output the footage with the V-709 color overlaid on the signal, the camera doesn't provide any way to use alternative LUTs. If you're using an external monitor/recorder, those monitors frequently have the ability to apply LUTs and some even allow loading in user-created LUTs. If you know the LUT you're going to be using in post, loading that into your monitor can make for a much more pleasant viewing experience during production. You would then configure the SYSTEM SETTINGS>COLOR SETTINGS>SDI OUT and/or HDMI OUT menu options to output the unmodified V-Log signal, and that signal is suitable for recording on the external recorder, and you can then choose to have the external monitor overlay a viewing LUT to make it look more like what the finished footage may look like.

You can (and should) use the zebras and the waveform monitor and the spotmeter as exposure tools for V-Log, but you have to use them differently than you would for a video gamma. The levels that you may be familiar with (such as 70% IRE for skin tones, and 100% IRE for highlights) simply aren't appropriate for use with V-Log; you'll have to adapt and learn to use new levels when judging the exposure of your footage.

There are two general schools of thought when exposing raw footage or logarithmic gamma footage: exposing for middle gray, or Expose To The Right (ETTR). Let's discuss ETTR first.

### Option 1: Exposing To The Right (ETTR)

Exposing To The Right is a technique based on using a histogram for exposure. A histogram shows the distribution of brightness in the image, and the further the image is shifted towards the highlights, the further right the image moves on the histogram. Proponents of the "Expose To The Right" technique argue that the darker tones and shadows are the noisiest parts of the image, so if you can lift your image up out of the shadow area, you can take advantage of the cleaner upper range of the sensor and gamma curve; later, you can drop the footage's levels back down to where proper exposure would dictate it should be. Furthermore, those who shoot raw footage frequently embrace ETTR because raw sensor data is generally stored linearly, not logarithmically; that means that the vast majority of available "bits" are allocated to the brightest stops, and the darkest stops receive the fewest "bits" (i.e., the least number of shades of gray that they can represent). Generally, the very brightest stops are allotted the most data and retain the finest gradations, more so than the darker stops do. This is indeed a concern with many cameras' raw footage, but is not nearly as much

of a concern when using a logarithmic gamma; a LOG gamma redistributes the bits so that they are more equally dispersed along the entire dynamic range of the sensor. It is true that the darkest stops will still have fewer “bits” allocated to them than the brightest stops will, but over much of the range they will be treated more or less equally, and as such, the value of lifting the shadows and dark tones up is less important when recording with a logarithmic curve than it would be with linear sensor data. Both the EVA1’s V-Log recording and raw output encode the signal logarithmically, so this is generally not a concern on the EVA1.

The general idea behind ETTR is to expose the image as bright as you possibly can, so long as none of the video information “clips” off the top. Regardless of how dark an image should look in the final footage, the idea is that if you expose it in the top part of the exposure range, you’ll get the cleanest, lowest-noise images, and you can always push it back down to proper exposure in post. ETTR proponents use the histogram as an exposure tool to accomplish this, because the histogram plots out all the exposure levels in a given scene, and if there is any unused space on the right edge of the histogram, that means that you have room to brighten up the exposure (which will shift the histogram’s graph over to the right within its frame; hence Expose To The Right). In theory, it sounds great; on still photographs, it can work great.

Now, as a general technique there’s nothing wrong with ETTR, it does work and it is a reasonable choice. However, it is not necessarily the best choice, because ETTR is designed to preserve the highlights with no consideration of what happens to the midtones. It can result in retaining a lot of detail and in making the least-noisy footage, but it will mandate extensive post-production correction on every single shot. When exposing using ETTR, skin tones may end up being recorded brighter or darker in every scene, simply based on where the highlights happen to be in that particular shot, and every shot will need to be corrected to bring the skin tones back to a reasonably consistent level so that your footage will intercut cleanly and seamlessly. And, depending on just how bright the highlights are in any given scene, ETTR may result in a scenario where the skin tones and midtones are significantly underexposed in an effort to catch and preserve all the highlights. That might make for nice highlights but it might also result in noisy skin tones and midtones, since in post production you may have to stretch the skin tones up out of the darker (and noisier) sections of the sensor. Generally, cinematography is (and should be) more about the

subject than it should be about the highlights; excessive attention to the highlights may mean compromising other aspects of the footage, so a strict “ETTR” approach is not always going to provide the overall best results in a video project.

The ETTR approach has traditionally relied on using a histogram to set exposure. The AU-EVA1 doesn’t offer a histogram, it offers a waveform monitor. Still, the same principle applies -- the exposure would generally be set to saturate the waveform monitor and fill it up top to bottom. You could perhaps say that you’re using “ETTT”, as in, Expose To The Top, when using the waveform. You’ll still want to keep the brightest areas under control though, or you risk clipping one or more color channels in the very brightest sections of the footage, which can cause color shifts in your highlights.

## Option 2: Exposing For Middle Gray

An alternative method of exposure would be to expose for middle gray. Video systems are frequently referencing “middle gray” or “18% gray.” 18% gray is a photographic and film standard, it’s a shade of gray that reflects 18% of the light that hits it. It is frequently incorporated into test charts, and you can easily buy an “18% gray card” at photography stores. 18% reflectiveness represents approximately the average overall brightness of many scenes, and camera autoexposure systems are typically designed to expose to where the scene represents approximately 18% reflectance levels. In Ansel Adams’ “Zone System”, middle gray is known as Zone V.

When exposing for middle gray, you’ll find the zebras and the waveform monitor extremely useful. In this section I’ll refer to exposure levels in terms of IRE units. In conventional video gammas, middle gray is usually exposed properly at somewhere around 50 to 55 IRE. However, not so in V-Log. In V-Log, middle gray is properly represented at 42 IRE. If you happen to have an 18% gray card in your scene, it should show up on your waveform monitor at approximately 42 IRE for “proper” exposure (note, here I am using the term “proper” in an idealized, mathematical way; the artistic merits of the scene may very well dictate that the exposure needs to be higher or lower than this).

The V-Log gamma curve maps the following brightness levels to the following IRE levels:

0% reflectance (black):	7.3 IRE
18% reflectance (middle gray):	42 IRE
90% reflectance:(white):	61 IRE
absolute clipped superwhite:	109 IRE

If you are using test charts, you will likely have access to 18% “middle gray” and 90% “white”; many gray cards sold in photographic stores will have 18% gray on one side, and 90% white on the other. 90% reflectance doesn’t necessarily indicate “pure white” or the brightest object that can be seen or recorded; rather it is (as its name suggests) a white where 90% of the light that hits it is reflected. The camera is capable of seeing and rendering brightness above 90% reflectance, as illustrated by the fact that 90% reflectance is mapped to 61 IRE, and the camera can continue rendering brighter detail all the way to 109 IRE.

In V-Log on the EVA1, the curve is laid out so that there are 8 stops below middle gray, and 6 stops above middle gray. You can, of course, choose to modify that by underexposing middle gray some; if you underexpose by one stop, you’ll then have 7 stops below middle gray and 7 stops above it. In all cases you’ll get 14 stops of dynamic range; the recommended allocation is for middle gray to be at 42 IRE with 8 stops below and 6 stops above, but you can shift that on an as-needed basis, so long as you account for it in post. This is one reason why it is such an excellent idea to shoot a standardized test chart at the head of every scene, so the colorist knows exactly what the intended exposure was and can account for any individualized decisions that were made on a scene-to-scene basis.

The technique of exposing towards middle gray is similar to conventional video gamma exposure, where you frequently will have some typical “anchor points” in your exposure plan (such as having Caucasian skin highlights peak at about 70 IRE on a standard video gamma, and keeping your highlights at or below 100 IRE). Keeping skin tones comparable shot-to-shot makes matching footage easier in post, obviously; keeping middle gray levels constant will make matching V-Log footage easier in post too. (Note: I’m not saying that you should expose skin tones at middle gray level, nor am I saying that you may necessarily want to expose middle gray at 42 IRE; I’m just pointing out that exposing at a consistent level makes matching footage in post much easier.)

Exposing for a logarithmic gamma isn’t necessarily as simple as putting an 18% gray card in the scene and exposing it for 42 IRE. It can be that

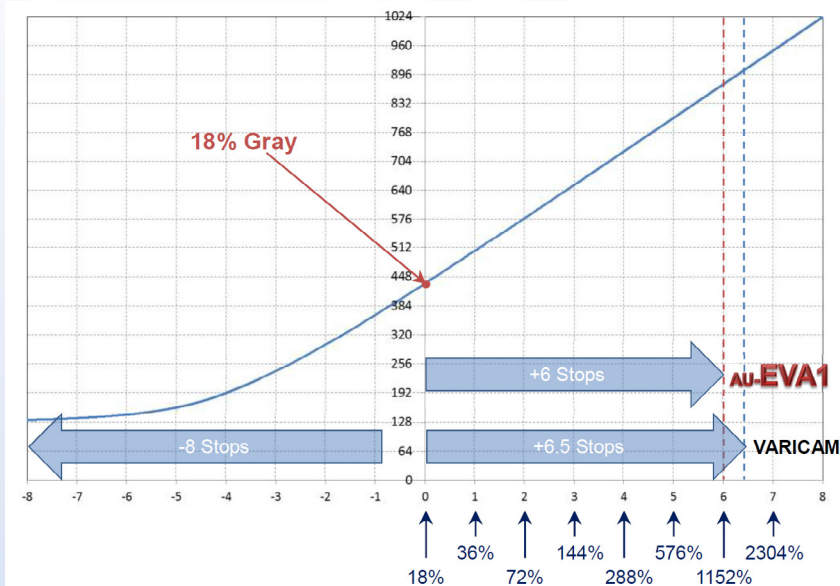


simple, if you want it to be, but there are steps you can take to perhaps improve the images the camera generates. The question is usually one of balancing noise versus retaining highlights. As with all digital camera sensors, the darkest regions of the image are typically the areas that show the most noise. Exposing To The Right is a technique designed to lift the image up out of the noisy area and have it render in the “sweet spot” of the sensor’s exposure range, which you can then drop down to proper exposure in post-production while avoiding some of the sensor noise. And that’s a valid technique, but it does sacrifice some of the sensor’s dynamic range (dropping the footage back down crushes off the darkest tones), and it can result in highly inconsistent midtones from shot to shot, which will require extensive correction in post. When you overexpose the image, you also run the risk of clipping the highlights earlier. When you properly expose the image, you may maintain the highlights but you may also encounter some noise in the shadows. If you are filming a scene where there are a lot of very bright highlights, you may actually need to underexpose the scene to preserve those highlights, even if it means pushing your subject down into the noisier darker sections of the sensor’s range. Or, you may just have to bite the bullet and accept that sometimes highlights clip and there’s nothing you can do about it — or, rather, there may be nothing that you should do about it; compromising the quality of the main subject in a quest for preserving highlights may not be an acceptable tradeoff in some cases. Of course the ideal solution is to adjust the lighting in the scene so that the highlights aren’t too bright; depending on the type of production you’re doing that may or may not be an option.

There is no one overall “right” answer, there is only a question of your priorities — if you cannot stand clipped highlights under any circumstances, maybe you should use ETTR, understanding that in the quest to preserve every bit of highlight detail you might end up underexposing the image, resulting in noisier images. If you cannot abide noise at all, maybe you should consider establishing a “noise floor,” an IRE level which you will not allow the important elements of your image to fall into. Perhaps you find the range from, say, 0 to 10 IRE too noisy for your tastes, so you may choose to overexpose your images so that the darkest significant details in the image are at least 10 IRE or brighter. Such overexposure may possibly result in clipped highlights, and will certainly lead to reduced dynamic range, but perhaps that’s the tradeoff you’re willing to make for minimizing any appearance of image noise.

## Adapting the Zebras And Waveform to V-Log

Whichever method you choose is a matter of personal preference. I find that the easiest way to work with V-Log is to expose for middle gray. Placing middle gray at 42 IRE, 90% white at 61 IRE, and black at 7 IRE gives a wide exposure range that allows for 6 stops of exposure over middle gray, and 8 stops under middle gray. Using these general exposure levels, you'll find that properly-exposed highlights on skin tones will usually fall between about 42 IRE for dark-skinned subjects, on up to about 55 IRE for light-skinned subjects. The full range of skin tones will of course be much wider; in this section I'm making some recommendations on what you might want to consider for the peak brightness levels on the skin tones (example: on a light-skinned Caucasian model, you might want to aim for the brightest portion of their skin — such as a reflection off the cheekbone or forehead — to peak at a maximum of about 55 IRE). Doing so will mean that the rest of the model's skin will fall in lower IRE ranges, but that the peak will be controlled at no higher than 55 IRE. It also means that shot-to-shot, the density of tonality should be similar between shots without the need for substantial correction. When shooting in REC 709 gamma, I'm used to setting my zebras at 70 and 100; I would use 70 to check highlights on skin, aiming for there to be just a little bit of 70-IRE area on Caucasian skin tones highlights; then I would check Zebra 2 to check for highlight clipping. With Zebra 2 set at 100 IRE, that would show areas that were approaching clipping, without necessarily meaning that they were indeed clipping. Modifying this approach for V-Log, I now set Zebra 1 at 55 IRE, and Zebra 2 at 90 IRE. In this way I can use the same monitoring tools that I'm used to, in the same way I'm used to, but gain the wider latitude and post flexibility of V-Log.



In the above graph, Panasonic has charted the full V-Log curve, for both the EVA1 and for the Varicam. Middle gray is represented by the 18% mark in the center of the chart. The legend across the bottom represents the number of f-stops away from middle gray for each vertical column; you can see that the chart accounts for 8 stops below middle gray and up to 8 stops above middle gray. V-Log was designed to account for sensors that deliver up to 16 f-stops of dynamic range. But the EVA1's sensor can't generate that much; the EVA1's sensor is capable of 14 stops, and the Varicam's sensor is capable of about 14.5 stops.

If you decide to underexpose to retain more highlights, be aware that you're going to have to stretch the shadows and midtones up more in post, and that will likely reveal more noise that will have to be dealt with by noise reduction in post.

### A Word About Recording 10-bit vs. 8-bit

The EVA1's internal recordings can be done at 8-bit quantizing and 4:2:0 color sampling, or 10-bit quantizing and 4:2:2 color sampling. AVCHD always records as 8-bit 4:2:0, and the ALL-I modes always record as 10-bit 4:2:2. Some of the LongGOP recordings are 8-bit 4:2:0 and some are 10-bit 4:2:2. 8-bit 4:2:0 recordings may be perfectly suitable for many purposes and jobs, but 8-bit is not as robust as 10-bit, and 4:2:0 is not as robust as 4:2:2 color sampling (obviously).

When working with V-Log, you'll be stretching and pushing the footage quite a bit in post. You will probably experience notably better results with a 10-bit 4:2:2 recording, than you would from an 8-bit 4:2:0 recording. It's not that you can't work with 8-bit 4:2:0, it's just that -- well, 10-bit 4:2:2 is better. The further you push the footage, the more 10-bit and 4:2:2 will hold up as compared to 8-bit and 4:2:0. When recording V-Log, I recommend you record in the highest-bitrate, highest-color sampling codec you can.