



Data Levels, Color Management, ACES, and HDR

6

Data Levels, Color Management, ACES, and HDR

This chapter covers operational details that affect how color is managed for media that is imported into and exported from DaVinci Resolve. If color accuracy is important to you, then it's a good idea to learn more about how Resolve handles the data levels of each clip, how DaVinci Resolve Color Management helps you to work with different formats, and how to use ACES.

Data Levels Settings and Conversions	134
Internal Image Processing and Clip Data Levels	134
Assigning Clip Levels in the Media Pool	135
Video Monitoring Data Levels	135
Deck Capture and Playback Data Level	136
Output Data Level Settings in the Deliver Page	136
So, What's the "Proper" Data Range for Output?	137
DaVinci Resolve Color Management	137
Display Referred vs. Scene Referred Color Management	137
Resolve Color Management for Editors	138
Resolve Color Management and Camera Raw Formats	138
Controlling the Input, Timeline, and Output Color Space	139
Procedures for Using Resolve Color Management	141
Exporting Color Space Information to QuickTime Files	142
Color Management Using ACES	142
Setting Up ACES in the Project Settings Window	143
Tips For Rendering Out of an ACES Project	145
High Dynamic Range (HDR) Grading in Resolve	145
The Different Formats of HDR	146
Dolby Vision	147
SMPTE ST.2084, Ultra HD Premium, and HDR 10	154
Simultaneous SDR and HDR Output While Grading	158

Data Levels Settings and Conversions

Different media formats use different ranges of values to represent image data. Since these data formats often correspond to different output workflows (cinema vs. broadcast), it helps to know where your project's media files are coming from, and where they're going, in order to define the various data range settings in DaVinci Resolve and preserve your program's data integrity.

To generalize, with 10-bit image values (with a numeric range of 0-1023), there are two different data levels (or ranges) that can be used to store image data when writing to media file formats such as QuickTime, MXF, or DPX. These ranges are:

- **Video:** Typically used by Y'CBCR video data. All image data from 0 to 100 percent must fit into the numeric range of 64–940. Specifically, the Y' component's range is 64–940, while the numeric range of the CB and CR components is 64–960. The lower range of 4–63 is reserved for “blacker-than-black,” and the higher ranges of 941/961–1019 are reserved for “super-white.” These “out of bounds” ranges are recorded in source media as undershoots and overshoots, but they're not acceptable for broadcast output.
- **Full:** Typical for RGB 444 data acquired from digital cinema cameras, or film scanned to DPX image sequences. All image data from 0 to 100 percent is simply fit into the full numeric range of 4 to 1023.

Keep in mind that every digital image, no matter what its format, has absolute minimum and maximum levels, referred to in this section as 0–100 percent. Whenever media using one data range is converted into another data range, each color component's minimum and maximum data levels are remapped so that the old minimum value is scaled to the new data level minimum, and the old maximum value is scaled to the new data level maximum:


- (minimum Video Level) 64 = 4 (Data Level minimum)
- (maximum Video Level) 940 or 960 = 1023 (Data Level maximum)

Converting an image from one data range to another should result in a seamless change. All “legal” data from 0–100 percent is always preserved, and is linearly scaled from the previous data range to fit into the new data range.

The sole exception is that the undershoots and overshoots of “Video Levels” media, if present, are usually clipped when converted to full-range “Full Levels.” However, Resolve preserves this data internally, and these clipped pixels of detail in the undershoots and overshoots are still retrievable by making suitable adjustments in the Color page to bring them back into the “legal” range.

Internal Image Processing and Clip Data Levels

It's useful to know that, internally to DaVinci Resolve, all image data is processed as full range, uncompressed, 32-bit floating point data. What this means is that each clip in the Media Pool, whatever its original bit-depth or data range, is scaled into full-range 32-bit data. How each clip is scaled depends on its Levels setting in the Clip Attributes window, available from the Media Pool contextual menu.



Selecting Video or Full levels

By converting all clips to uncompressed, full-range, 32-bit floating point data, Resolve guarantees the highest quality image processing that's possible. As always, the quality of your output is dependent on the quality of the source media you're using, but you can be sure that Resolve is preserving all the data that was present in the original media.

Assigning Clip Levels in the Media Pool

When you first import media into the Media Pool, either manually in the Media page or automatically by importing an AAF or XML project in the Edit page, Resolve automatically assigns the "Auto" Levels setting. When a clip is set to Auto, the Levels setting used is determined based on the codec of the source media.

DaVinci Resolve generally does a good job of figuring out the appropriate Levels setting of each clip on its own. However, in certain circumstances, such as when you're mixing video-originated media and digital cinema camera-originated media, you may find that you need to manually choose the appropriate settings so that the levels of each clip are interpreted correctly. This can be done using each clip's Levels setting in the Clip Attributes window, available from the Media Pool contextual menu in either the Media or Edit pages.

To change a clip's Data Level setting:

- 1 Open the Media or Edit page.
- 2 Select one or more clips, then right-click one of them and choose Clip Attributes.
- 3 Click the Levels ratio button corresponding to the data level setting you want to assign, then click OK.

TIP: If you need to change the Levels setting of a range of clips that share a unique property such as reel name, resolution, frame rate, or file path, you can view the Media Pool by column, and sort by the particular column that will best isolate the range of media to which you need to make a data level assignment

Once you change a clip's Levels setting, that clip will automatically be reconverted based on the new assignment. If it appears to be correct, then you're ready to work. If it doesn't, then you may want to reconsider the Levels assignment you've made, and you should check with the person who provided the media to find out how it was generated, captured, and exported.

So long as the Levels settings used by your clips are accurate, you should be ready to work. However, problems can still occur based on what external video hardware you're using with your workstation, and how you need to deliver the finished media to your client. For this reason, there are three additional data level settings that you can use to maintain data integrity, while at the same time seeing the proper image as you work.

Video Monitoring Data Levels


Superficial problems may result if the settings used by your external display differ from the settings you're using to process data levels in Resolve. Accordingly, there is a Video/Full Level setting in the Master Project Settings panel of the Project Settings (in the Video Monitoring section).

When you change this setting, the image being output to your external display should change, but the image you see in your Viewer will not. That's because this setting only affects the data levels being output via the video interface connecting the Resolve workstation to your external display. It has no effect on the data that's processed internally by Resolve, or on the files written when you render in the Deliver page.

There are two options:

- **Video:** This is the correct option to use when using a broadcast display set to the Rec. 709 video standard. (10-bit 64-940).
- **Full:** If your monitor or projector is capable of displaying “full range” video signals, and you wish to monitor the full 10-bit data range (4-1023) while you work, then this is the correct option to use.

It is imperative that the option you choose in Resolve matches the data range the external display is set to. Otherwise, the video signal will appear to be incorrect, even though the internal data is being processed accurately by DaVinci Resolve.



Video/Full Level selection for monitoring

Deck Capture and Playback Data Level

There is a separate “Video/Data Level” setting that is specific to when you’re capturing from or outputting to VTRs. This setting also affects the video signal that is output via the video interface connecting the Resolve workstation to your VTR (which is usually also in the signal chain used for monitoring). However, it only takes effect when you’re capturing from tape in the Media page, or editing to tape in the Deliver page. If you never capture or output to tape, this setting will never take effect.

This setting is found in the Deck Capture and Playback panel of the Project Settings.

The reason for a separate option for tape capture and output is that often you’d want to monitor in one format (normally scaled Rec. 709), but output to tape in another (full range RGB 444). This way, you can set up Resolve to accommodate this workflow, and then not have to worry about manually switching your video interface back and forth.

There are two options:

- **Video:** This is the correct option to use when you want to output conventional Rec. 709 video to a compatible tape format.
- **Full:** This is the correct option to use when you want to output “full range” RGB 444 video to a compatible tape format.

Once tape ingest or output has finished, your video interface goes back to outputting using the setting specified by the “Colorspace conversion uses” setting in the Master Project Settings panel of the Project Settings (in the Video Monitoring section).

Output Data Level Settings in the Deliver Page

Finally, there’s one last set of data level settings, available in the Render Settings list, within the Format group. It’s the “Set to video or data level” pop-up menu. It’s there to give you the ability to convert the data level of your rendered output, if necessary.

All media is output using a single data level, depending on your selection. There are three options:

- **Automatic:** The output data level of all clips is set automatically based on the codec you select to render to in the “Render to” pop-up menu.
- **Video:** All clips are rendered as normally scaled for video (10-bit 64-940).
- **Full:** All clips are rendered as full range. (10-bit 4-1019)

For most projects, leaving this setting on “Automatic” will yield the appropriate results. However, if you’re rendering media for use by another image processing application (such as a compositing application) that is capable of handling “full range” data, then full range output is preferable for media exchange as it provides the greatest data fidelity. For example, when outputting media for VFX work as a DPX image sequence, or as a ProRes 4444 encoded QuickTime file, choosing “Unscaled full range data” guarantees the maximum available image quality. However, it is essential that the application you use to process this media is set to read it as “full range” data, otherwise the images will not look correct.

So, What’s the “Proper” Data Range for Output?

Strictly speaking, there is no absolutely “proper” data range to use when outputting image data. As long as the Levels setting of each clip in the Media Pool is set to reflect how each clip was created, your primary consideration is which data range is compatible with the media format or application you’re delivering to. If the media format you’re exporting to supports either normally scaled or full range, and the application that media will be imported into supports either normally scaled or full range, then it’s really your choice, as long as everyone involved with the project understands how the data range of the media is meant to be interpreted once they receive it.

Outputting to hardware is a bit trickier, in that you need to make sure that the external display or VTR you’re outputting to is set up to receive a signal using the data range you’ve chosen. If the device is limited to only one data range, then you need to be sure that you’re outputting to it using that data range, or the levels of the image will appear to be incorrect, even though the image data being processed by Resolve is actually fine.

DaVinci Resolve Color Management

How color is managed in DaVinci Resolve depends on the “Color Science” setting in the Project Settings. There are four options: DaVinci YRGB, DaVinci YRGB Color Managed, DaVinci ACES, and DaVinci ACESc. This section discusses the second setting, DaVinci YRGB Color Managed. ACES is discussed in the next section in this chapter.

Display Referred vs. Scene Referred Color Management

The default DaVinci YRGB color science setting, which is what Resolve has always used, relies on what is called “Display Referred” color management. This means that Resolve has no information about how the source media used in the Timeline is supposed to look; you can only judge color accuracy via the calibrated broadcast display you’re outputting to. Essentially, *you* are the color management, in conjunction with a trustworthy broadcast display that’s been calibrated to ensure accuracy.

Starting in DaVinci Resolve 12, there’s a new color science option called “DaVinci YRGB Color Managed,” or more simply “Resolve Color Management” (RCM). This is a so-called “Scene Referred” color management scheme, in which you have the option of matching each type of media you’ve imported into your project with a color profile that informs DaVinci Resolve how to represent each specific color from each clip’s native color space into the common working color space of the timeline in which you’re editing, grading, and finishing.

This is important, because two clips that contain the same RGB value for a given pixel may in actuality be representing different colors at that pixel, depending on the color space that was originally associated with each captured clip. This is the case when you compare raw clips shot with different cameras made by different manufacturers, and it’s especially true if you compare clips recorded using the differing log-encoded color spaces that are unique to each camera.

Scene Referred color management via RCM doesn't do your grading for you, but it does try to ensure that the color and contrast from each different media format you've imported into your project are represented accurately in your timeline. For example, if you use two different manufacturer's cameras to shoot green trees, recording Blackmagic Film color space on one, and recording to the Sony SGamut3.Cine/SLog3 color space on the other, you can now use RCM to make sure that the green of the trees in one set of clips match the green of the trees in the other, within the shared color space of the Timeline.

This sort of thing can also be done manually in a more conventional Display Referred workflow, by assigning LUTs that are specific to each type of media, in order to transform each clip from the source color space to the destination color space that you require. However, RCM uses mathematical transforms, rather than lookup tables, which makes it easier to extract high-precision, wide-latitude image data from each supported camera format, in order to preserve high-quality image data from acquisition, through editing, color grading, and output. RCM is also easier to use, freeing you from the need to locate and maintain a large number of LUTs to accommodate your various workflows.

The preservation of wide-latitude image data deserves elaboration. LUTs clip image detail that goes outside of the numeric range they're designed to handle, so this often requires the colorist to make a pre-LUT adjustment to "pull back" image data in the highlights that you want to retrieve. Using RCM eliminates this two-step process, since the input color space math used to transform the source preserves all wide-latitude image data, making highlights easily retrievable without any extra steps.

How is DaVinci Resolve Color Management Different from ACES?

This is a common question, but the answer is pretty simple. Resolve Color Management (RCM) and ACES are both Scene Referred color management schemes designed to solve the same problem. However, if you're not in a specific ACES-driven cinema workflow, DaVinci Resolve Color Management can be simpler to use, and will give you all of the benefits of color management, while preserving the the same "feel" that the DaVinci Resolve Color page controls have always had.

Resolve Color Management for Editors

RCM is also easier for editors to use in situations where the source material is log-encoded. Log-encoded media preserves highlight and shadow detail, which is great for grading and finishing, but it looks flat and unpleasant, which is terrible for editing.

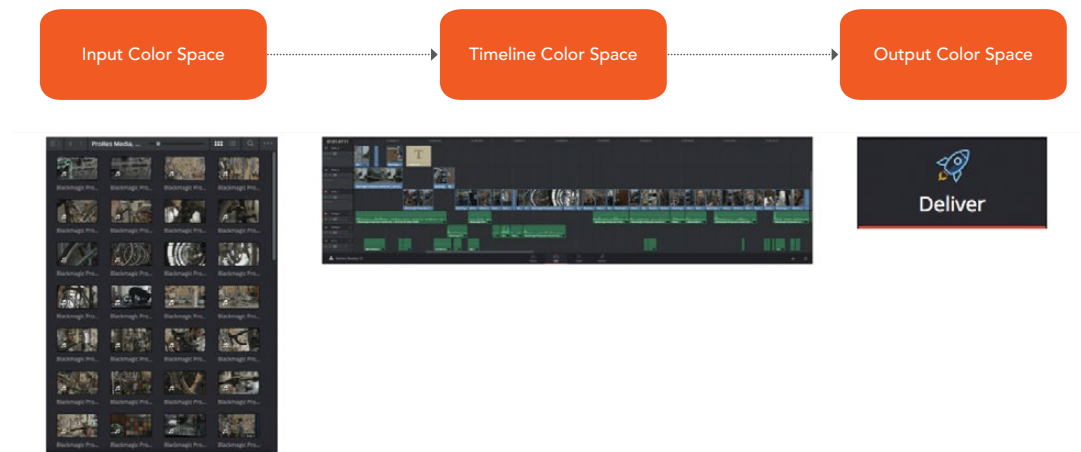
Even if you have no idea how to do color correction, it's simple to turn RCM on in the Project Settings, and then use the Media Pool to assign the particular Input Color Space that corresponds to the source clips from each camera. Once that's done, each log-encoded clip is automatically normalized to the default Timeline Color Space of Rec.709 Gamma 2.4. So, without even having to open the Color page, editors can be working with pleasantly normalized clips in the Edit page.

Resolve Color Management and Camera Raw Formats

When you use RCM in a project that uses Camera Raw formats, color science data from each camera manufacturer is used to debayer each camera raw file to specific color primaries with linear gamma, so that all image data from the source is preserved and made available to Resolve's color managed image processing pipeline. As a result, the Camera Raw project settings and Camera Raw palette of the Color Page are disabled, because RCM is controlling the debayering of all camera raw clips, and all image data from the raw file is available no matter which Timeline Color Space you choose to work within.

Controlling the Input, Timeline, and Output Color Space

An additional benefit of using Scene Referred color management via RCM is that, not only do you have the ability to identify the color science of each source media format (the Input Color Space), but you also have explicit control over the working color space (the Timeline Color Space), and you have separate control over the Output Color Space. Basically, Resolve Color Management consists of two color transforms working together, converting the source clips via assignable Input Color Spaces to the Timeline Color Space in which you work, and then converting the Timeline Color Space to whatever Output Color Space you require to deliver the project.



Resolve Color Management consists of three color transforms working together

This means that, as a colorist, you can set the Timeline Color Space that you're working in to whatever you prefer. If you prefer grading log media because you like the way the grading controls behave in that color space, you can set the Timeline Color Space in the Color Management panel of the Project Settings to any of the available log formats, including ARRI Log C and Cineon Film Log. If you instead prefer grading in the Rec. 709 color space because you're more comfortable with how the controls feel in that color space, you can choose that instead. Whatever Timeline Color Space you assign is what all source clips will be transformed to for purposes of making grading adjustments in the Color page, so you can make this choice using a single setting.

A key benefit of the color space conversions that RCM applies is that no image data is ever clipped during the Input to Timeline color space conversion. For example, even if your source is log-encoded or in a camera raw format, grading with a Rec. 709 Timeline Color Space does nothing to clip or otherwise limit the image data available to the RCM image processing pipeline. All image values greater than 1.0 or less than 0.0 are preserved and made available to the next stage of RCM processing, the Timeline to Output color space conversion.


Consequently, if you're grading in a color space other than the one you need to output to, you don't have to worry about data loss during the color transformation back to the color space you actually want to output to. The Output Color Space setting gives you the freedom to work using whatever Timeline Color Space you like while grading, with Resolve automatically converting your output to the specific color space you want to monitor with and deliver to. And thanks to the precision of the image processing in DaVinci Resolve, you can convert from a larger color space to a smaller one and back again without clipping or a loss of quality. Of course, if you apply a LUT or use Soft Clip within a grade, then clipping will occur, but that's a consequence of using those particular operations.

TIP: If you want to use Resolve Color Management, but you want the Input and Output Color Spaces to match whatever you set the Timeline Color Space to, you can choose “Bypass” in the “Input Colorspace” and “Output Colorspace” pop-up menus.

Finally, it is the Output Color Space that determines the final color space of your rendered result. While no image data is clipped during the Source to Timeline color space conversion, image data will be clipped during the Timeline to Output color space conversion in order for the final image to conform to the color space being rendered and output.


Single Setting vs. Dual Setting RCM

Starting in Resolve 12.5, there are two ways you can set up RCM. When the “Use Separate Color Space and Gamma” checkbox is turned off, the Color Management panel of the Project Settings exposes one pop-up each for the Input, Timeline, and Output Color Space settings. Each setting lets you simultaneously transform the gamut and gamma, depending on which option you choose. This makes it a bit simpler to set up the transform you need.




Single setting Resolve Color Management

If you turn the “Use Separate Color Space and Gamma” checkbox on, then the Color Management panel changes so that the Input, Timeline, and Output Color Space settings each display two pop-ups. The first pop-up lets you explicitly set the gamut, while the second pop-up lets you explicitly set the gamma. This makes it easier to see exactly which pair of transforms is being used at each stage of RCM.



Dual setting Resolve Color Management

Additionally, Dual Setting RCM enables you to assign separate gamut and gamma transforms to clips in the Media Pool.



Dual setting Resolve Color Management assignments for Media Pool clips

Procedures for Using Resolve Color Management

Despite the seeming complexity of color management, using it is actually pretty simple, since the goal is to enable Resolve to automatically take care of as much of the complexity as possible, freeing you up to focus on doing creative work. In essence, all you have to do is (A) turn RCM on, (B) assign the appropriate Input Color Space to each clip in the Media Pool based on its source, and (C) choose the Timeline and Output Color Space combination you want to use.

To enable Resolve Color Management:

Open the Master Project Settings panel of the Project Settings, and choose DaVinci YRGB Color Managed from the “Color Science” pop-up menu.

To choose the default color space for all clips with an unassigned Input Color Space:

Open the Color Management panel of the Project Settings, choose whether you want to use Single Setting or Dual Setting color management, and then choose the desired options from the “Input Colorspace” pop-up menu or menus. This setting determines the color space of all clips you haven’t manually assigned an Input Color Space to. The default setting is “Rec.709 Gamma 2.4.” If you choose “Bypass,” then all clips will default to whatever color space is selected in the “Timeline Colorspace” pop-up menu.

To assign an Input Color Space to one or more selected clips in the Media Pool:

- 1 Select the clip or clips you want to assign. Since you can assign an Input Color Space to more than one clip at a time, it will speed things up if you use a sort or find operation, or create a Smart Bin or Smart Filter, to identify every clip from a specific camera using a particular format, so you can assign them all at once.
- 2 Do one of the following:
 - In the Media Pool, right-click one of the selected clips, and choose the Input Color Space and/or Input Gamma (if Dual Setting color management is enabled) that corresponds to those clips from the contextual menu.
 - In the Color page, right-click any clip’s thumbnail in the Thumbnail timeline, and choose the Input Color Space and/or Input Gamma (if Dual Setting color management is enabled) that corresponds to those clips from the contextual menu.

To change the Timeline and Output Color Space:

Open the Color Management panel of the Project Settings, and choose the desired options from the single menu or pair of pop-up menus for the “Timeline Colorspace” and “Output Colorspace” RCM settings. The default settings are “Rec.709 Gamma 2.4.” If you choose “Bypass” in the “Output Colorspace” pop-up menu, then the Output Color Space will match whatever you’ve selected in the “Timeline Colorspace” pop-up menu.

All in all, using RCM is an easy way to ensure high image quality and a predictable starting point when working on projects requiring log-encoded conversion, or that incorporate multiple media formats.

Exporting Color Space Information to QuickTime Files

If you render QuickTime files from the Deliver page, then color space tags will be embedded into each file based on either the Timeline Color Space (if Resolve Color Management is disabled) or the Output Color Space (if Resolve Color Management is enabled). The following tags will be written, if currently selected:

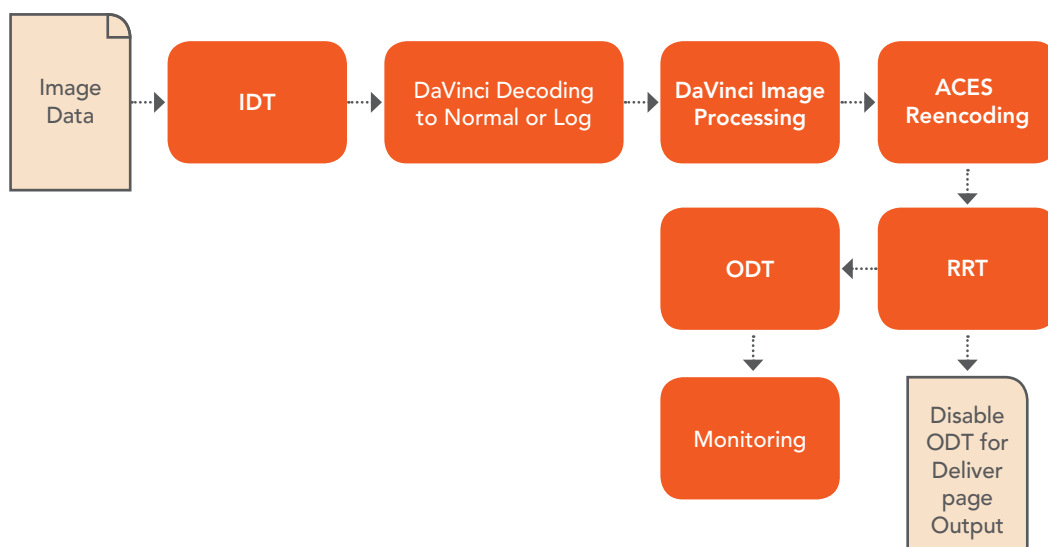
- REC 709
- REC 2020
- CIE XYZ
- P3DCI
- P3D60

Color Management Using ACES

The ACES (Academy Color Encoding Specification) color space has been designed to make scene-referred color management a reality for high-end digital cinema workflows. ACES also makes it easier to extract high-precision, wide-latitude image data from raw camera formats, in order to preserve high-quality image data from acquisition through the color grading process, and to output high-quality data for broadcast viewing, film printing, or digital cinema encoding.

An oversimplification of the way ACES works is that every camera and acquisition device is characterized to create an IDT (Input Device Transform) that specifies how media from that device is converted into the ACES color space. The ACES gamut has been designed to be large enough to encompass all visible light, with more than 25 stops of exposure latitude. In this way ACES has been designed to be future-proof, taking into consideration advances in image capture and distribution.

Meanwhile, an RRT (Reference Rendering Transform) is used to transform the data provided by each image format's IDT into standardized, high-precision, wide-latitude image data that in turn is processed via an ODT (Output Device Transform). Different ODT settings correspond to each standard of monitoring and output, and describe how to accurately convert the data within the ACES color space into the gamut of that display in order to most accurately represent the image in every situation. The RRT and ODT always work together.



ACES signal and processing flow

By using the ACES color space and specifying an IDT and an ODT, you can ingest media from any capture device, grade it using a calibrated display, output it to any destination, and preserve the color fidelity of the graded image.

Setting Up ACES in the Project Settings Window

There are three parameters available in the Project Settings that let you set up DaVinci Resolve to use the ACES workflow. The first is located in the Project Settings panel:

- **Color science is:** Using this pop-up menu, you can choose either DaVinci ACES, or DaVinci ACEScc color science, which enables ACES processing throughout DaVinci Resolve.
- **DaVinci ACES:** Chooses DaVinci ACES color science, which applies a proprietary encoding to the ACES data before it is processed by Resolve. After processing, a reverse encoding is applied in order to output ACES linear data.
- **DaVinci ACEScc:** Choose DaVinci ACEScc color science to apply a standard Cineon-style log encoding to the ACES data before it is processed by Resolve. This well defined common encoding makes it possible for ASC CDL values to be used across systems using the same ACEScc encoding. After processing, a reverse encoding is applied in order to output ACES linear data.
- **ACES Version:** When you've chosen one of the ACES color science options, this pop-up becomes available to let you choose which version of ACES you want to use. While DaVinci Resolve incorporates ACES 1.0, pre-release versions of ACES are still available for projects requiring backward compatibility, including versions 0.1, 0.2, and 0.7.

When you choose the ACES color science, two additional settings become available in the Color Management panel:

- **ACES IDT:** This pop-up menu lets you choose which IDT (Input Device Transform) to use for the dominant media format in use. DaVinci Resolve currently supports the following IDTs:
 - **ALEXA/BMD Film/4K/4.6K/Canon 1D/5D/7D/C300/C500/Panasonic V35/Sony RAW/slog2/slog3:** Each listed camera has its own transform in this list that is specific to each camera's sensor. Some of these transforms (specifically the Canon C300 and C500 transforms) have multiple selections for daylight and tungsten, and for different combinations of gamut and gamma.
 - **Rec. 709 (Camera):** A deprecated legacy IDT for Rec. 709 that's included for backward compatibility. Converts the source data to linear based on Rec. 709 and transforms the result to ACES, but while this transformation is technically correct, it's not necessarily pleasing after conversion through the matching ODT. For this reason, the academy updated to the following Rec. 709 IDT, which is the inverse of the Rec. 709 ODT.
 - **Rec. 709:** A standard transform designed to move media in the Rec. 709 color space into the ACES color space. This option is used for any other file type that might be imported, such as ProRes from Final Cut Pro, DNxHD from Media Composer, and any media file captured from tape.
 - **Rec. 709 (D60 Sim):** A standard transform designed to move media in the Rec. 709 color space with a white point of D60 into the ACES color space.
 - **Rec. 2020 (Camera):** A deprecated legacy IDT for Rec. 2020 that's included for backward compatibility.
 - **Rec. 2020:** This IDT transforms media created with the wide-gamut standard for consumer and broadcast television.

- **DCDM (Camera):** A deprecated legacy IDT for DCDM that's included for backward compatibility.
- **DCDM:** This IDT transforms X'Y'Z'-encoded media with a gamma of 2.6.
- **P3 D60:** Transforms RGB-encoded image data with a D60 white point, intended for monitoring with a P3-compatible display using a D60 white point.
- **P3 D60 (Camera):** A deprecated legacy IDT for P3 D60 that's included for backward compatibility.
- **P3 D60 PQ (1000/2000/4000 nits):** Transforms an image that's compatible with the P3 color gamut, using the SMPTE standard PQ (ST.2084) tone curve for High Dynamic Range (HDR) postproduction. Three settings for three different peak luminance ranges are provided; which one is appropriate to use depends on the maximum white level of the display used to create the media. Preliminary standards exist for HDR displays with peak luminance at 1000 nits, 2000 nits, and 4000 nits.
- **P3 DCI (Camera):** A deprecated legacy IDT for P3 DCI that's included for backward compatibility.
- **ADX (10 or 16):** 10-bit or 16-bit integer film-density encoding transforms meant for use if you're working with film scans that were initially encoded in an ACES workflow. This transform is designed to maintain the variation in look between different film stocks.
- **sRGB (D60 Sim.):** A standardized transform designed for media created for computer display in a consumer environment.

If you're working on a project that mixes media formats that require different IDTs, then you can assign different IDTs to clips using the Media Pool's contextual menu, or using the Clip Attributes window, which is also accessible via the Media Pool's contextual menu.

- **ACES ODT:** This pop-up menu lets you choose an ODT (Output Device Transform) with which to transform the image data for monitoring on your calibrated display, and when exporting a timeline in the Deliver page. You can choose from the following options:
 - **Rec. 709:** This ODT is used for standard monitoring and deliverables for TV.
 - **Rec. 2020:** This ODT is for compatibility with the wide-gamut standard for consumer and broadcast television.
 - **DCDM:** This ODT exports X'Y'Z'-encoded media with a gamma of 2.6 intended for handoff to applications that will be re-encoding this data to create a DCP (Digital Cinema Package) for digital cinema distribution. This can be displayed via an XYZ-capable projector.
 - **P3 D60:** Outputs RGB-encoded image data with a D60 white point, intended for monitoring with a P3-compatible display using a D60 white point.
 - **P3 D60 PQ (1000/2000/4000 nits):** Outputs an image that's compatible with the P3 color gamut, using the SMPTE standard PQ tone curve for High Dynamic Range (HDR) postproduction. Three settings for three different peak luminance ranges are provided; which one is appropriate to use depends on the maximum white level of your display. Preliminary standards exist for HDR displays with peak luminance at 1000 nits, 2000 nits, and 4000 nits.
 - **P3 DCI:** Outputs RGB-encoded P3 image data with the native P3 white point, intended for monitoring with a P3-compatible display (using its native white point).
 - **ADX (10 and 16):** A standardized ODT designed for media destined for film output. Two settings accommodate 10-bit and 16-bit output. This ODT is not meant to be used for monitoring.

- **sRGB (D60 Sim.):** A standardized ODT designed for media destined for computer display in a consumer environment. Suitable for monitoring when grading programs destined for the web.
- **P3 DCI:** A standardized ODT that outputs RGB-encoded image data with a D61 white point, intended for use when outputting media for DCI mastering.

You must manually select an ODT that matches your workflow and room setup when working in ACES.

The Initial State of Clips When Working in ACES

Don't worry if the initial state of each image file appears differently than what was monitored originally on set. What's important is that if the camera original media was well exposed, the IDT used in ACES mode will retain the maximum amount of image data, and provide the maximum available latitude for grading, regardless of how the image initially appears on the Timeline.

Tips For Rendering Out of an ACES Project

When choosing an output format in the Deliver page, keep the following in mind:

- If you've delivering graded media for broadcast, set the ACES Output Device Transform to be Rec. 709, then you can output to whatever media format is convenient for your workflow.
- When you're delivering graded media files to another ACES-capable facility using the DCDM or ADX ODCs, you should choose the OpenEXR RGB Half (uncompressed) format in the Render Settings, and set the ACES Output Device Transform to "No Output Device Transform."
- When you're rendering media for long-term archival, you should choose the OpenEXR RGB Half (uncompressed) format in the Render Settings, and set the ACES Output Device Transform to "No Output Device Transform."

High Dynamic Range (HDR) Grading in Resolve

The HDR features found in DaVinci Resolve are only available in DaVinci Resolve Studio.

High Dynamic Range (HDR) video describes an emerging family of video encoding and distribution technologies designed to enable a new generation of television displays to play video capable of intensely bright highlights and increased saturation. The general idea is that the majority of an HDR image will be graded similarly to how a Standard Dynamic Range (SDR) image is graded now, with the shadows, midtones, and low highlights being mostly the same between traditionally SDR and HDR-graded images, to maintain a comfortable viewing experience and easier backward compatibility. However, HDR provides abundant additional headroom for very bright highlights and color saturation that far exceed what has been visible in SDR television and cinema, without clipping, that enable the colorist to create more vivid and life-like highlights in images such as sunsets, lit clouds, firelight, explosions, sparkles, and other intensely bright and colorful imagery. This not only provides more life-like lighting intensity and saturation, but it also dramatically expands the contrast available in the scene. For example, an SDR display may have a peak luminance level of 100 "nits" (cd/m²), but HDR displays may have peak luminance levels of 500, 1000, or even 4000 nits.

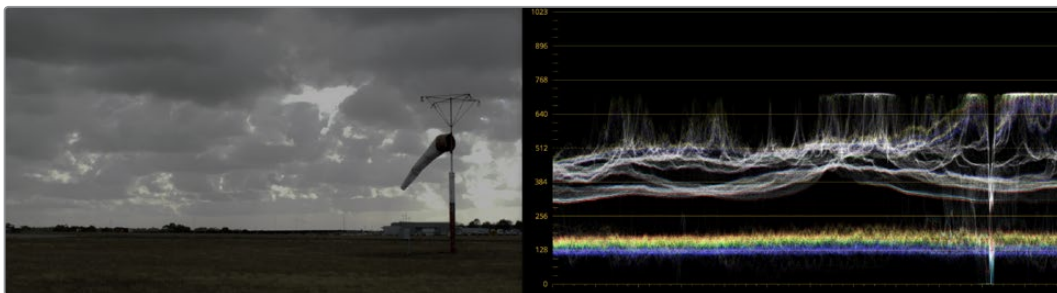
However, because it's an emerging technology, the technical standards being proposed far exceed what the first few generations of consumer displays are capable of. At the time of this writing, consumer televisions are only capable of outputting closer to 300, 500, or 800 nits, and are saddled with automatic brightness limiting (ABL) circuits in order to limit power consumption to acceptable levels for home use, which means that only a certain percentage of the picture may reach these peak values at any one time. This is fine, because the point of HDR is not that you're making the entire image brighter, it's that you have more headroom for high saturation and specific bright highlights.

For this reason, HDR standards focus on describing what displays should be capable of, not how these levels are to be used. This means that the question of how to utilize the expansive headroom for brightness and saturation that HDR enables is fully within the domain of the colorist, as a series of creative decisions that must be made regarding how to assign the varying levels of brightness that are available in your source media to the HDR standard you're mastering to as you grade.

The Different Formats of HDR

While different HDR technologies use different methods to map the video levels of your program to an HDR display's capabilities, they all output a "near-logarithmically" encoded signal that requires a compatible television that's capable of correctly stretching this signal into its "normalized" form for viewing. This means if you look at an HDR signal that's output from the video interface of your grading workstation on an SDR display, it will look flat, desaturated, and unappealing until it's plugged into your HDR display of choice.

At the time of this writing, there are three approaches to mastering HDR that DaVinci Resolve is capable of supporting, including Dolby Vision, HDR 10 and ST.2084 (derived from the PQ curve developed by Dolby), and Hybrid Log-Gamma (HLG). Each of these HDR standards focuses on describing how an HDR signal is encoded for output and later mapped to the output of an HDR display.



A graded HDR image being output looks similar to a log-encoded image

Each of these standards is most easily enabled using Resolve Color Management (RCM) via Color Space options in the Color Management panel of the Project Settings. Alternately, LUTs are available for each of these color space conversions if you want to do things the old fashioned way, but Resolve Color Management has become so mature in the last year that, from experience, I personally recommend this approach to handling HDR in Resolve.

However, these standards have nothing to say about how these HDR-strength levels are to be used creatively. This means that the question of how to utilize the expansive headroom for brightness and saturation that HDR enables is fully within the domain of the colorist, as a series of creative decisions that must be made regarding how to assign the range of highlights that are available in your source media to the above-100 nit HDR levels you're mastering to as you grade, given the peak luminance level that you're mastering with. Which HDR peak luminance level you use (1000 nit, 4000 nit) probably depends on which monitor you have access to, and which HDR mastering technology you're using.

The following sections will describe how to work with Dolby Vision, HDR 10, and Hybrid Log-Gamma in Resolve.

Dolby Vision

Long a pioneer and champion of the concept of HDR for enhancing the consumer video experience, Dolby Labs has developed a method for delivering HDR called Dolby Vision. Dolby Vision uses the SMPTE 2084 “PQ” EOTF along with colorist-controlled metadata that maintains a program’s artistic intent across multiple displays capable of a wide luminance range, from 0 to 10,000 cd/m². Resolve provides access to the SMPTE 2084 PQ EOTF settings required for mastering Dolby Vision, both as Color Space settings in RCM and as a set of 3D LUTs. While Dolby Vision content is not limited to a particular color space, Resolve Color Management provides a P3 D65 setting that matches the capabilities of the current generation of mastering monitors available from Dolby.

However, to accommodate backward compatibility with SDR displays, as well as the varying maximum brightness of different makes and models of HDR consumer displays, Dolby Vision offers both dual layer backward compatible and single layer non-backward compatible output options (mastering is identical for both).

The dual layer backward compatible output option accommodates a two-stream video delivery method that consists of a base layer and an enhancement layer with metadata. On an SDR television, only the base layer is played, which contains a Rec. 709-compatible image that’s a colorist-guided approximation of the HDR image. On a Dolby Vision-enabled HDR television, however, both the base and enhancement layers will be recombined, using additional “artistic guidance” metadata generated by the colorist to determine how the resulting HDR image highlights should be scaled to fit the varied peak white levels and color volume performance that’s available on any given Dolby Vision television.

Those, in a nutshell, are the twin advantages of the Dolby Vision system. It’s backward compatible with SDR televisions, and it’s capable of intelligently scaling the HDR highlights to provide the best representation of the mastered image for whatever peak luminance and color volume a particular television is capable of. All of this is guided by decisions made by the colorist during the grade.

So, who’s using Dolby Vision? At the time of this writing, all seven major Hollywood studios are mastering in Dolby Vision for Cinema. Studios that have pledged support to master content in Dolby Vision for home distribution include Universal, Warner Brothers, Sony Pictures, and MGM. Content providers that have agreed to distribute streaming Dolby Vision content include Netflix, Vudu, and Amazon. If you want to watch Dolby Vision content on television at home, consumer display manufacturers LG, TCL, Vizio, and HiSense have all announced models with Dolby Vision support.

DaVinci Resolve Hardware Setup for Dolby Vision

To master with Dolby Vision in DaVinci Resolve, you need a somewhat elaborate hardware setup, consisting of the following equipment:

- Your DaVinci Resolve grading workstation, outputting via either a DeckLink 4K Extreme 12G or an UltraStudio 4K Extreme video interface
- A Dolby Vision Certified HDR Mastering Monitor
- An SDR (probably Rec. 709-calibrated) display
- A standalone hardware video processor called the Content Management Unit (CMU), which is a standard computer platform with a Video I/O card. The CMU is only available from Dolby Authorized System Integrators; please contact Dolby for an Authorized Systems Integrator near you.
- A video router such as the BMD Smart Videohub

This hardware is all connected as seen in the following illustration:

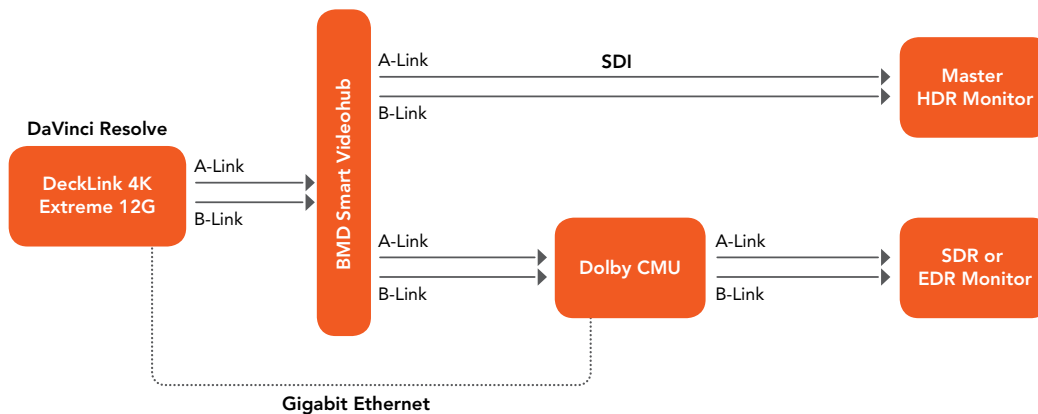


Illustration of how to connect the necessary equipment for Dolby Vision grading in Resolve

Your Resolve workstation will connect its dual SDI outputs to the BMD Smart Videohub, which splits the video signal to two mirrored sets of SDI outputs. One mirrored pair of SDI outputs goes to your HDR display. The other mirrored pair of SDI outputs goes to the CMU (Content Mapping Unit), which is itself connected to your SDR display via SDI. Lastly, the Resolve workstation is connected to the Dolby CMU via Gigabit Ethernet to enable the CMU to communicate back to Resolve.

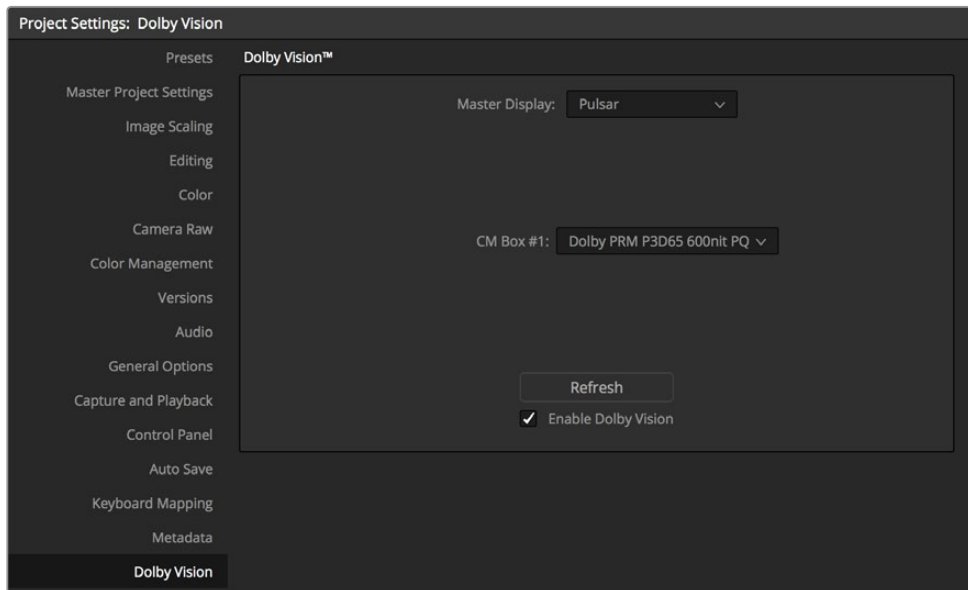
The CMU is an off-the-shelf video processor that uses a combination of proprietary automatic algorithms and colorist-adjustable metadata within Resolve to define how an HDR-graded video should be transformed into an SDR picture that can be displayed on a standard Rec. 709 display, as well as how the enhancement layer should scale itself to varying peak luminance levels.

Authorizing Resolve to Work With Dolby Vision

To expose the Dolby Vision controls in DaVinci Resolve Studio, you need a Dolby Vision Mastering license from Dolby. To install this license, refer to the instructions provided with the license from Dolby. After installation, analyze any clip using the Color > Dolby Vision submenu controls, and the Dolby Vision palette should become visible.

Dolby Vision Controls in DaVinci Resolve


When your DaVinci Resolve workstation is authorized, a variety of controls become available in DaVinci Resolve, starting with the Dolby Vision panel in the Project Settings, which exposes settings for choosing what kind of Master Display you're using, how the CMU should be configured to output, a button to let you refresh your connection to the CMU, and a checkbox that enables and disables Dolby Vision controls.



The Dolby Vision panel that appears in the Project Settings


All of the Dolby Vision analysis commands appear in the Color > Dolby Vision submenu. There are four commands that let you analyze clips in the Timeline in different ways:

- **Analyze All Shots:** Automatically analyzes each clip in the Timeline and stores the results individually.
- **Analyze Selected Shot(s):** Only analyzes selected shots in the Timeline.
- **Analyze Selected And Blend:** Analyzed multiple selected shots and averages the result, which is saved to each clip. Useful to save time when analyzing multiple clips that have identical content.
- **Analyze Current Frame:** A fast way to analyze clips where a single frame is representative of the entire shot.



The Analyze commands in the Color > Dolby Vision menu

A Dolby Vision palette exposes a set of six controls, consisting of luminance-only Lift/Gamma/Gain controls (that work slightly differently than those found in the Color Wheels palette), Saturation, Chroma Weight Offset (which darkens parts of the picture to preserve colorfulness that's clipping in Rec. 709), and Tone Detail Weight, which preserves contrast detail in the highlights that might otherwise be lost when the highlights are mapped to lower dynamic ranges, usually due to clipping. Currently, this control is disabled for 100 nit/Rec. 709 grades (for grades higher than 100 nits, increasing Tone Detail Weight increases the amount of highlight detail that's preserved).



Controls found in the Dolby Vision palette are only visible once you've authorized your system with a special key, available from Dolby

These Dolby Vision analysis and trim controls in DaVinci Resolve send metadata to the CMU by encoding it into the first line of the SDI output. This metadata guides how the CMU makes this transformation, because the CMU is actually the functional equivalent of the Dolby Vision chip that's inside each Dolby Vision-enabled television; what you're really doing is using the CMU to make your SDR display simulate a 100 nit Dolby Vision television.

Additionally, the CMU can be used to output 600 nit, 1000 nit, and 2000 nit versions of your program, if you want to see how your master will scale to those peak luminance levels. This, of course, requires the CMU to be connected to a display that's capable of being set to those peak luminance output levels.

Though not required, you have the option to visually trim your grade at up to four different peak luminance levels, including 100 nit, 600 nit, 1000 nit and 2000 nit reference points, so you can optimize a program's visuals for the peak luminance and color volume performance of many different televisions with a much finer degree of control. If you take this extra step, Dolby Vision-compatible televisions will use the artistic guidance metadata you generate in each trim pass to ensure the creative intent is preserved as closely as possible, in an attempt to provide the viewer with the best possible representation of the director's intent.

For example, if a program were graded relative to a 4000 nit display, along with a single 100 nit Rec. 709 trim pass, then a Dolby Vision-compatible television with 750 nit peak output will reference the 100 nit trim pass artistic guidance metadata in order to come up with the best way of "splitting the difference" to output the signal correctly. On the other hand, were the colorist to do three trim passes, the first at 100 nits, a second at 600 nits, and a third at 1000 nits, then a 750 nit-capable Dolby Vision television would be able to use the 600 and 1000 nit artistic intent metadata to output more accurate HDR-strength highlights that take better advantage of the 750 nit output of that television.

Dolby Vision Certified Mastering Monitors

At the time of this writing, only three displays have been certified as Dolby Vision Certified Mastering Monitors. Requirements include a minimum peak brightness of 1000 nits, a 200,000:1 contrast ratio, greater than Rec. 709 color gamut (preferably P3), and native support for SMPTE ST.2084 as the EOTF (otherwise known as the PQ curve). These displays include:

- The Sony BVM X 300 (30", 1000 nit, 4K)
- The Dolby PRM 32FHD (32", 2000 nit, 1080)
- The Dolby Pulsar (42", 4000 nit, 1080)

Of these, only the Sony is commercially available. The Dolby monitors are not commercially available and are provided only in limited availability from Dolby.

When grading Dolby Vision, the HDR monitor should set to use P3/D65.

Setting Up Resolve Color Management For Grading HDR

Once the hardware is set up, setting up Resolve itself to output HDR for Dolby Vision mastering is easy using Resolve Color Management (RCM). This procedure is pretty much the same no matter which HDR mastering technology you're using; only specific Output Color Space settings will differ.

- 1 Set Color Science to DaVinci YRGB Color Managed in the Master Project Settings, and Option-click the Save button to apply the change without closing the Project Settings.
- 2 Then, open the Color Management panel, and set the Output Color Space pop-up to the HDR ST.2084 setting that corresponds to the peak luminance, in nits, of the grading display you're using. For example, if you're grading with a Sony BVM X300, choose HDR ST.2084 1000 nits. Be aware that whichever HDR setting you choose will impose a hard clip at the maximum nit value supported by that setting. This is to prevent accidentally overdriving HDR displays, which can have negative consequences.
 - HDR ST.2084 300 nits
 - HDR ST.2084 500 nits
 - HDR ST.2084 800 nits
 - HDR ST.2084 1000 nits
 - HDR ST.2084 2000 nits
 - HDR ST.2084 4000 nits

This setting is only the output EOTF (a sort of gamma transform, if you will, using the terminology that DaVinci Resolve's UI has used up until now).

- 3 Next, choose a setting in the Timeline Color Space that corresponds to the gamut you want to use for grading, and that will be output. For example, if you want to grade the Timeline as a log-encoded signal and "normalize" it yourself, you can choose ARRI Log C or Cineon Film Log. If you would rather have Resolve normalize the Timeline to P3-D65 and grade that way, you could choose that setting as well. In terms of defining the output gamut, the rule is that if "Use Separate Color Space and Gamma" is turned off, the Timeline Color Space setting will define your output gamut. If "Use Separate Color Space and Gamma" is turned on, then you can specify whatever gamut you want in the left Output Color Space pop-up menu, and choose the EOTF from the right pop-up menu (as described in step 2).
- 4 Be aware that, when it's being properly output, HDR ST.2084 signals appear to be very log-like, in order to pack their wide dynamic range into the bandwidth of a standard video signal. It's the HDR display itself that "normalizes" this log-encoded image to look as it should. For this reason, the image you see in your Color page Viewer is

going to appear flat and log-like, even though the image being displayed on your HDR reference display looks vivid and correct. If you want to make the image in the Color Page Viewer look “normalized” at the expense of clipping the HDR highlights, you can use the 3D Color Viewer Lookup Table setting in the Color Management panel of the Project Settings to assign the appropriate “HDR X nits to Gamma 2.4 LUT,” with X being the peak nit level of the HDR display you’re using.

- 5 Additionally, the “Timeline resolution” and “Pixel aspect ratio” (in the project settings) that your project is set to use is saved to the Dolby Vision metadata, so make sure your project is set to the final Timeline resolution and PAR before you begin grading.

Resolve Grading Workflow For Dolby Vision

Once the hardware and software is all set up, you’re ready to begin grading Dolby Vision HDR. The workflow is fairly straightforward.

- 1 First, grade the HDR image on your Dolby Vision Certified Mastering Monitor to look as you want it to. Dolby recommends starting by setting the look of the HDR image, to set the overall intention for the grade.
- 2 When using various grading controls in the Color page to grade HDR images, you may find it useful to enable the HDR Mode of the node you’re working on by right-clicking that node in the Node Editor and choosing HDR mode from the contextual menu. This setting adapts that node’s controls to work within an expanded HDR range. Practically speaking, this makes controls that operate by letting you make adjustments at different tonal ranges, such as Custom Curves, Soft Clip, etcetera, work more easily with wide-latitude signals.
- 3 When you’re happy with the HDR grade, click the Analysis button in the Dolby Vision palette. This analyzes every pixel of every frame of the current shot, and performs and stores a statistical analysis that is sent to the CMU to guide its automatic conversion of the HDR signal to an SDR signal.
- 4 If you’re not happy with the automatic conversion, use the Lift/Gamma/Gain/Chroma Weight/Chroma Gain controls in the Dolby Vision palette to manually “trim” the result to the best possible Rec. 709 approximation of the HDR grade you created in step 1. This stores what Dolby refers to as “artistic guidance” metadata.
- 5 If you obtain a good result, then move on to the next shot and continue work. If you cannot obtain a good result, and worry that you may have gone too far with your HDR grade to derive an acceptable SDR downconvert, you can always trim the HDR grade a bit, and then re-trim the SDR grade to try and achieve a better downconversion. Dolby recommends that if you make significant changes to the HDR master, particularly if you modify the blacks or the peak highlights, you should re-analyze the scene. However, if you only make small changes, then reanalyzing is not strictly required

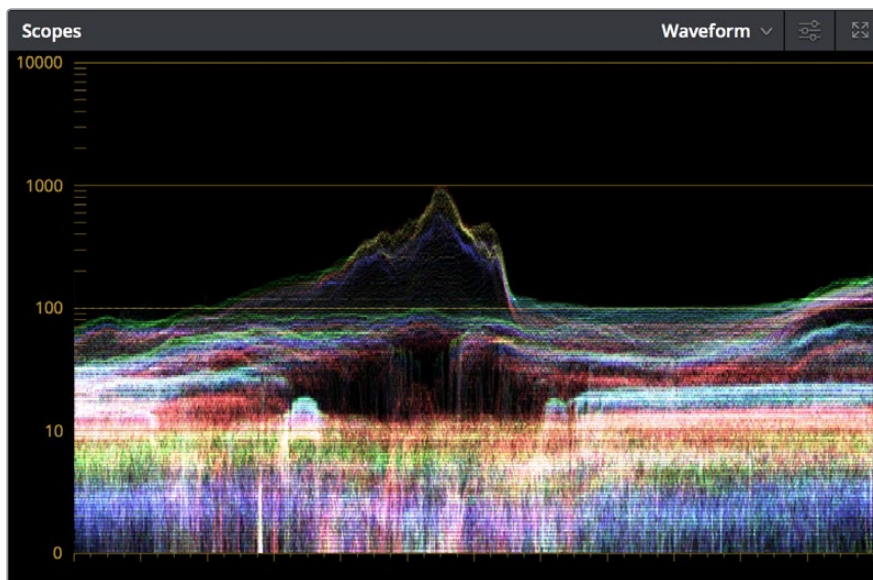
As you can see, the general idea promoted by Dolby is that a colorist will focus on grading the HDR picture relative to the 1000, 2000, 4000, or higher nit display that is being used, and will then rely on the colorist to use the Dolby Vision controls to “trim” this into a 100 nit SDR version with artistic guidance. This “artistic guidance” metadata is saved as part of the mastered media, and it’s used to more intelligently scale the HDR highlights to fit within any given HDR display’s peak highlights, to handle how to downconvert the image for SDR displays, and also how to respond when a television’s ABL circuit kicks in. In all of these cases, the colorist’s artistic intent is used to guide all dynamic adjustments to the content.

Analyzing HDR Signals Using Scopes

When you're using waveform scopes of any kind, including parade and overlay scopes, the signal will fit within the 10-bit scale much differently owing to the way HDR is encoded. The following chart of values will make it easier to understand what you're seeing:

- 1023 = 10,000 nits (no known display)
- 920 = 4000 nits (peak luminance on a Dolby Pulsar Monitor)
- 844 = 2000 nits (peak luminance on a Dolby PRM 32FHD)
- 767 = 1000 nits (peak luminance on a Sony BVM X300)
- 528 = 108 nits (Dolby Cinema projector peak luminance)
- 519 = 100 nits
- 447 = 48 nits (DCI projection and Dolby Cinema 3D peak luminance)
- 0 = 0 nits

If you're monitoring with the built-in video scopes in DaVinci Resolve Studio, you can turn on the "Enable HDR Scopes for ST.2084" checkbox in the Color panel of the Project Settings, which will replace the 10-bit scale of the video scopes with a scale based on "nit" values (or cd/m²) instead.



The Video Scopes with "Enable HDR Scopes for ST.2084" enabled in the Color panel of the Project Settings

If you're unsatisfied with the amount of detail you're seeing in the 0 – 519 range (0 – 100 nits) of the video scope graphs, then you can use the 3D Scopes Lookup Table setting in the Color Management panel of the Project Settings to assign the appropriate "HDR X nits to Gamma 2.4 LUT," with X being the peak nit level of the HDR display you're using. This converts the way the scopes are drawn so that the 0 – 100 nit range of the signal takes up the entire range of the scopes, from 0 through 1023. This will push the HDR-strength highlights up past the top of the visible area of the scopes, making them invisible, but it will make it easier to see detail in the midtones of the image.

Rendering a Dolby Vision Master

To deliver a Dolby Vision master after you've finished grading, you want make sure that the Output Color Space of the Color Management panel of the Project Settings is set to the appropriate HDR ST.2084 setting, based on the peak luminance in nits of your HDR display. Then, you want to set your render up to use one of the following Format/Codec combinations:

- TIFF, RGB 16-bit
- EXR, RBG-half (no compression)

When you render for tapeless delivery, the artistic intent metadata is rendered into a Dolby Vision XML and delivered with either the Tiffs or EXR renders. These two sets of files are then delivered to a facility that's capable of creating the Dolby Vision Mezzanine File (this cannot be done in Resolve).

Playing Dolby Vision at Home

On distribution, televisions that have licensed Dolby Vision use the base layer and enhancement layer+metadata to determine how the HDR image should be rendered given each display's particular peak luminance capabilities. Distributors, for their part, need to provide a minimum 10-bit signal to accommodate Dolby Vision's wide range. As a result, DolbyVision videos will look as they should on displays from 100 nits through 1000 nits peak luminance and above, scaling the enhancement layer's HDR-strength highlights to whatever peak luminance level is possible on that display and recombining these highlights with the base layer, with no unpredictable clipping.

SMPTE ST.2084, Ultra HD Premium, and HDR 10

Many display manufacturers who have no interest in licensing Dolby Vision for inclusion in their displays are instead going with the simpler method of engineering their displays to be compatible with SMPTE ST.2084. It requires only a single stream for distribution, there are no licensing fees, no special hardware is required to master for it (other than an HDR mastering display such as the Sony X300), and there's no special metadata to write or deal with (at this time).

Interestingly, SMPTE ST.2084 ratifies the "PQ" EOTF that was developed by Dolby and that's used by Dolby Vision that accommodates displays capable of peak luminance up to 10,000 cd/m² into a general standard. This standard requires at minimum a 10-bit signal for distribution, and the EOTF is described such that the video signal utilizes the available code values of a 10-bit signal as efficiently as possible, while allowing for such a wide range of luminance in the image.

SMPTE ST.2084 is also part of the new "Ultra HD Premium" industry specification, which stipulates that televisions bearing the Ultra HD Premium logo have the following capabilities:

- A minimum UHD resolution of 3840 x 2160
- A minimum gamut of 90% of P3
- A minimum dynamic range of either 0.05 nits black to 1000 nits peak luminance (to accommodate LCD displays), or 0.0005 nits black to 540 nits peak luminance (to accommodate OLED displays)
- Compatibility with SMPTE ST.2084

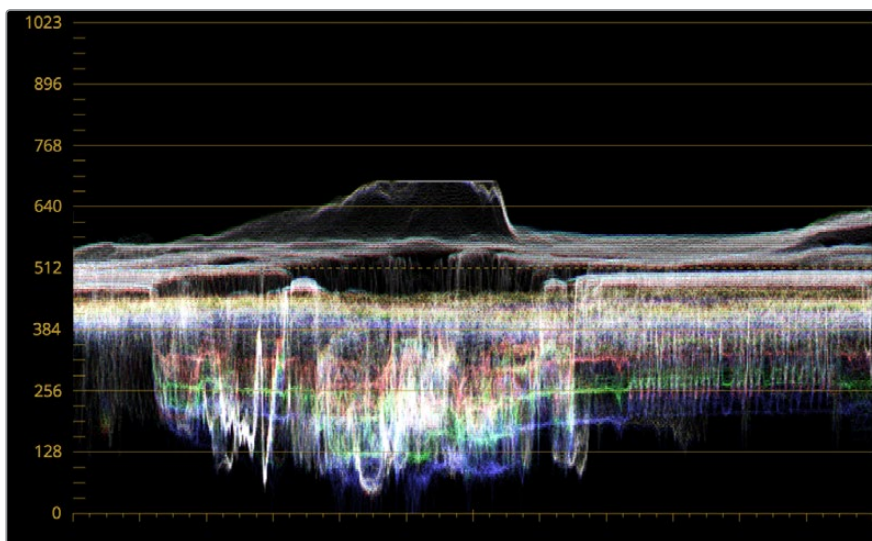
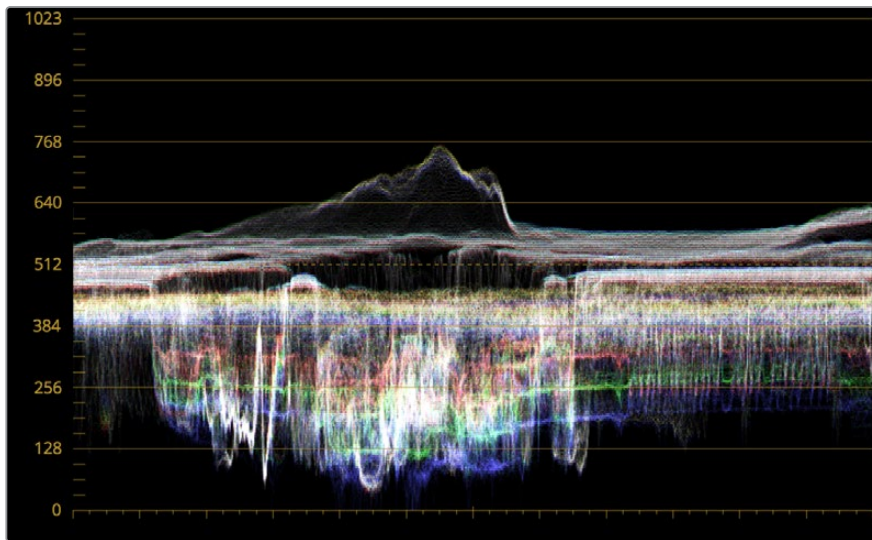
Finally, ST.2084 has been included in the HDR 10 standard adopted by the Blu-ray Disc Association (BDA) that covers Ultra HD Blu-ray. HDR 10 stipulates that Ultra HD Blu-ray discs have the following characteristics:

- UHD resolution of 3840 x 2160
- Up to the Rec. 2020 gamut

- SMPTE ST.2084
- Mastered with a peak luminance of 1000 nits

The downside is that, by itself, this EOTF is not backwards compatible with Rec. 709 displays using BT.1886 (although the emerging metadata standard SMPTE ST.2086 seeks to address this). Furthermore, no provision is made to scale the above-100 nit portion of the image to accommodate different displays with differing peak luminance levels. For example, if you grade and master an image to have peak luminance of 4000 nits, and you play that signal on an ST.2084-compatible television that's only capable of 800 nits, then everything above 800 nits will be clipped, while everything below 800 nits will look exactly as it should relative to your grade.

This is because ST.2084 is referenced to absolute luminance. If you grade an HDR image referencing a 1000 nit peak luminance display, as is recommended by HDR10, then any display using ST.2084 will respect and reproduce all levels from the HDR signal that it's capable of reproducing as you graded them, up to the maximum peak luminance level it can reproduce. For example, on an HDR10 compatible television capable of outputting 500 nits, all mastered levels from 501 – 1000 will be clipped, as seen in the screenshot below.



Comparing the original 1000 nit waveform representing the grading monitor to the 500 nit clipped waveform representing the consumer television

How much of a problem this is really depends on how you choose to grade your HDR-strength highlights. If you're only raising the most extreme peak highlights to maximum HDR-strength levels, then it's entirely possible that the audience might not notice that the display is only outputting 800 nits worth of signal and clipping any image details from 801 – 1000 nits because there weren't that many details above 800 anyway. Or, if you're grading large explosive fireballs up above 800 nits in their entirety because it looks cool, then maybe the audience will notice. The bottom line is, when you're grading for displays that are only capable of ST.2084, you need to think about these sorts of things.

Monitoring and Grading to ST.2084 in DaVinci Resolve


Monitoring an ST.2084 image is as simple as getting a ST.2084-compatible HDR display (such as the Sony X300), and connecting the output of your video interface to the input of the display. In the case of the Sony X300, which is a 4K capable display, you can connect four SDI outputs from a DeckLink 4K Extreme 12G with the optional DeckLink 4K Extreme 12G Quad SDI daughtercard, or an UltraStudio 4K Extreme, directly from your grading workstation to the X300, and you're ready to go.

Setting up Resolve Color Management to grade for ST.2084 is identical to setting up to grade for Dolby Vision. You'll also monitor the video scopes identically, and output a master identically, given that both standards rely upon the same PQ curve.

TIP: If you're monitoring with the built-in video scopes in Resolve, you can turn on the "Enable HDR Scopes for ST.2084" checkbox in the Color panel of the Project Settings, which will replace the 10-bit scale of the video scopes with a scale based on "nit" values (or cd/m²) instead.


Connecting to HDR-Capable Displays using HDMI 2.0a

If you have a DeckLink 4K Extreme 12G or an UltraStudio 4K Extreme video interface, then DaVinci Resolve 12.5 and above can output the metadata necessary to correctly display HDR video signals to display devices using HDMI 2.0a when you turn on the "Enable HDR metadata over HDMI" checkbox in the Master Project Settings.



The Enable HDR metadata over HDMI option in the Master Project Settings lets you output HDR via HDMI 2.0a

When you do so, a setting in the Color Management panel of the Project Settings, "HDR mastering is for X" lets you specify the output, in nits, to be inserted as metadata into the HDMI stream being output, so that the display you're connecting to correctly interprets it. The output you specify should match what your display is expecting.



The HDR mastering is for setting lets you insert metadata for HDR output via HDMI 2.0a

Hybrid Log-Gamma (HLG)

The BBC and NHK jointly developed another method of encoding HDR video, referred to as Hybrid Log-Gamma (HLG). The goal of HLG was to develop a method of mastering HDR video that would support a range of displays of different peak luminance capabilities without additional metadata, that could be broadcast via a single stream of data, that would fit into a 10-bit signal, and that in the words of the ITU-R Draft Recommendation BT.HDR, “offers a degree of compatibility with legacy displays by more closely matching the previous established television transfer curves.”

The basic idea is that the HLG EOTF functions very similarly to BT.1886 from 0 to 0.6 of the signal (with a typical 0 – 1 range), while 0.6 to 1.0 smoothly segues into logarithmic encoding for the highlights. This means that, if you just send an HDR Hybrid Log-Gamma signal to an SDR display, you’d be able to see much of the image identically to the way it would appear on an HDR display, and the highlights would be compressed to present an acceptable amount of detail for SDR broadcast.

On a Hybrid Log-Gamma compatible HDR display, however, the log-like highlights of the image (not the BT.1886-like bottom portion of the signal, just the highlights) would be stretched back out, relative to whatever peak luminance level a given HDR television is capable of outputting, to return the image to its true HDR glory. This is different from the HDR10 method of distribution described previously, in which the graded signal is referenced to absolute luminance levels dictated by ST.2084, and levels that cannot be represented by a given display will be clipped.

And while this facility to support multiple HDR displays with differing peak luminance levels is somewhat analogous to Dolby Vision’s ability to tailor HDR output to the unique peak luminance levels of any given Dolby Vision-compatible television, HLG requires no additional metadata to guide how the highlights are scaled, which depending on your point of view is either a benefit (less work), or a deficiency (no artistic guidance to make sure the highlights are being scaled in the best possible way).

As is true for most things, you don’t get something for nothing. The BBC White Paper WHP 309 states that, for a 2000 cd/m² HDR display with a black level of 0.01 cd/m², up to 17.6 stops of dynamic range without visible quantization artifacts (“banding”) is possible. BBC White Paper WHP 286 states that the proposed HLG EOTF should support displays up to about 5000 nits. So, partially, the backwards compatibility that HLG makes possible is due in part to discarding long-term support for 10,000 nit displays. However, given that the brightest commercially-available HDR display at the time of this writing is only 1000 nits peak luminance (the Sony X300), and the brightest HDR display currently in use outputs 4000 nits peak luminance (the experimental Dolby Pulsar), it’s an open question whether or not over 5000 nits is even necessary for consumer enjoyment.

At the time of this writing, only Sony has demonstrated consumer HDR televisions capable of displaying HLG encoded video. DaVinci Resolve, however, supports this standard through Resolve Color Management.

Grading Hybrid Log-Gamma in DaVinci Resolve

Monitoring an ST.2084 image is as simple as getting a Hybrid Log-Gamma-compatible HDR display, and connecting the output of your video interface to the input of the display.

Setting up Resolve Color Management to grade for HLG is identical to setting up to grade for Dolby Vision, except that there are two HLG settings to choose from for the Output Color Space:

- HDR HLG-2020
- HDR HLG-709

Optionally, if you choose to enable “Use Separate Color Space and Gamma,” you can choose either Rec. 2020 or Rec. 709 as your gamut, and HLG as your EOTF.

The levels you’ll be monitoring in your scopes will be different from the table of data to “nit” values listed previously for grading to the PQ EOTF.

Simultaneous SDR and HDR Output While Grading

It’s usually the case that you need to deliver both SDR and HDR versions of a program as deliverables. Opinions vary about whether it’s better to grade the HDR version first and then do a trim pass to produce the SDR version, versus grading the SDR version first and then doing a trim pass to produce the HDR version. Either way, it can be a big help to see both the HDR and SDR signals output at the same time to displays that are side-by-side, so you can directly compare them as you trim.

The Dolby Vision workflow accounts for this by virtue of the hardware setup that’s required, but if you’re grading using a more generic HDR 10 or HLG workflow, there’s still a way you can set this up using the Stereo 3D features that have been upgraded in DaVinci Resolve 12.5.1.

What you do is to convert your media and timeline to stereo, so that you can use the Stereo 3D palette to keep all of your SDR grades associated with one eye’s output, and all your HDR grades associated with the other eye’s output. You can switch back and forth between each eye using the Stereo 3D palette’s controls to choose which version you’re grading, and as you work, the SDR and HDR versions of each shot will be simultaneously output to different displays via your compatible Blackmagic Design video interface’s dual SDI outputs. It’s a bit of a hack, but it works.

You Can’t Use Resolve Color Management While Outputting SDR and HDR

Resolve Color Management (RCM) cannot be used in this workflow, because it doesn’t let you choose different Output Color Space settings for each SDI output. Instead, you must turn RCM off, and use the built-in HDR LUTs found in the 3D LUT submenu of each Corrector node’s contextual menu in the Node Editor of the Color page to convert one of the two streams into HDR.

The easiest way to do this is to apply the necessary LUT using the Timeline grade, since that’s the very last operation applied to every clip in a timeline. When you work on a timeline in stereo, the left and right eye streams of the Timeline grade are capable of holding separate corrections, which is a new feature as of 12.5.1.

Setting Up to Display Dual Video Streams via Two SDI Outputs

Originally intended for high-quality Stereo 3D monitoring, two separate SDI signals can be simultaneously output at full resolution using one of the following Blackmagic Design video interfaces:

- DeckLink HD Extreme 3D+
- DeckLink 4K Extreme
- DeckLink 4K Extreme 12G
- UltraStudio 4K
- UltraStudio 4K Extreme
- UltraStudio 4K Extreme 3

The following procedures describe how to set up stereo 3D monitoring in two different ways.

Monitoring via dual SDI to dual SDI:

- 1 Open the Master Project Settings panel of the Project Settings, then do the following:
 - Make sure Video Connection is set to YUV 4:2:2 SDI.
 - Turn on the “Use left and right eye SDI output” checkbox.
- 2 Open the Stereo 3D palette, and do the following:
 - Set Vision to Stereo.
 - Set the Out pop-up menu to None.

NOTE: When “Enable dual SDI 3D monitoring” is turned on, split-screen wipes and cursors will not be visible on the grading monitor, nor will you be able to view image resizing.

Converting Mono Clips or an Entire Timeline to Stereo

If you want to grade an HDR and non-HDR version of your timeline at the same time while outputting each to a separate display for monitoring, you need to do the following:

- Convert the clips used by that timeline from mono into stereo so each clip can have separate grades for each of the two SDI streams you’ll be outputting.
- Convert that timeline into a stereo timeline to enable dual-SDI output.
- Apply an HDR LUT to the Timeline grade of the left- or right-eye video stream so that stream can be output to the SDI output being connected to an HDR display.

First, you need to convert your media into stereo clips. This won’t actually do anything to add depth effects to the clips, it just flags them to be able to contain two grades capable of being output separately.

To convert one or more mono clips into stereo clips:

- 1 Select one or more non-stereo clips in the Media Pool. In this case, you’ll want to convert all of the clips used by a particular timeline you’re grading.
- 2 Right-click one of the selected clips and choose Convert to Stereo from the contextual menu.
 - Afterwards, the selected clips appear in the Media Pool as Stereo 3D clips and can expose their controls in the 3D Stereo palette of the Color page when they appear in a timeline.

Once you have a timeline full of clips that you've converted into stereo using the above procedure, you need to take the additional step of setting the Timeline to stereo in order to enable dual-SDI output of that timeline.

To convert a timeline to have to stereo for simultaneous HDR/SDR output while grading:

- Right-click a timeline in the Media Pool and choose Timelines > Set Timeline to Stereo

Once you've set a timeline and its media to be faux Stereo 3D, you're ready to set up which stream should be graded and output as HDR.

To set one of the stereo streams to output HDR:

- 1 Open the Stereo 3D palette, and click the Left or Right button to choose which video stream you want to monitor, grade, and output as HDR.
- 2 Turn off the Link button that's to the right of the Left and Right buttons to make sure that each stream's grade is unlinked. This is necessary to keep the SDR and HDR grades separate.
- 3 Choose the Timeline grade mode in the Node Editor, and press Option-S to create a new node.
- 4 Right-click the node you just created, and choose the desired output option from the 3D LUT > HDR ST.2084 or 3D LUT > HDR Hybrid Log-Gamma submenus. The output option you choose should correspond to the type of HDR display you have, and in the case of ST.2084, the maximum luminance value your display can output.

At this point, you're all set up to grade and output the SDR version of your program using one eye of the Stereo 3D palette, with the HDR version of your program graded and output via the other eye using the 3D LUT you've chosen.

To choose which stereo stream you want to grade:

- Open the Stereo 3D palette, and click the Left or Right button to choose which video stream you're grading.

The Left and Right buttons of the Stereo 3D palette choose which video stream you're set to grade in every single clip of the timeline. If you look at the small 3D badge at the bottom of each clip in the Thumbnail timeline, Orange indicates you're grading the left eye stream, and light blue indicates you're grading the right eye stream.

Outputting Both Video Streams in the Deliver Page

After you've graded both the SDR and HDR streams, you can choose to render one or both of them using the Render Stereoscopic 3D settings in the Video panel of the Render Settings in the Deliver page. You can choose to render "Both eyes as" with the option "Separate Files," or you can choose to render just the Left eye or Right eye media.