Affordable HDR Conversion with SL-HDR1 Metadata

A Cobalt Digital White Paper

By Ciro A. Noronha, Ph.D.

High Dynamic Range (HDR) is an emerging technology that delivers noticeably better image quality than standard dynamic range (SDR) content with minimal bandwidth increase. However, while HDR continues to gain popularity, SDR needs to be preserved – at least for the near future – because a large number of legacy consumer televisions do not support HDR. Acknowledging the mix of HDR and SDR devices in the marketplace, this white paper will examine different ways to provide cost-effective, simultaneous service to both types of devices, and explain why the SL-HDR1 conversion option is an ideal choice for broadcasters.

Luminance and Wide Color Gamut

Dynamic range is the ratio between the lowest and highest values of the luminance, or intensity of light emitted, from a display. Essentially, it is the ratio between the "whitest white" and "blackest black." Conventional 8-bit displays, for example, have a dynamic range of approximately 100:1. Dynamic range is usually expressed in stops, which is calculated as the log base 2 of the ratio. So, 8-bit SDR has a dynamic range of approximately six stops, while professional 10-bit SDR offers approximately 10 stops.

The human eye, however, can see about 12-14 stops, which means it can perceive more dynamic range than offered by 10-bit SDR material. So, how do you pack an HDR signal into a 10-bit display? Thankfully, the human eye is nonlinear in its response and perceives more detail at lower luminosities. HDR delivers images with improved details by assigning bits to light intensity in a nonlinear manner. More bits are assigned to the lower intensities and fewer bits to the higher intensities to express a higher range. Fundamentally, HDR shows more detail in the bright areas.

Light intensity is measured in candelas per square meter (cd/m2), also known as nits. A standard HDTV can produce luminance at about 100 nits. A UHD LCD display can range from 0.05 to 1,000 nits, while a UHD OLED monitor can produce 0.0005 to 540 nits. In contrast, HDR can code up to 10,000 nits, which no commercial monitor today can reproduce.

Total Luminance for SDR vs HDR Television

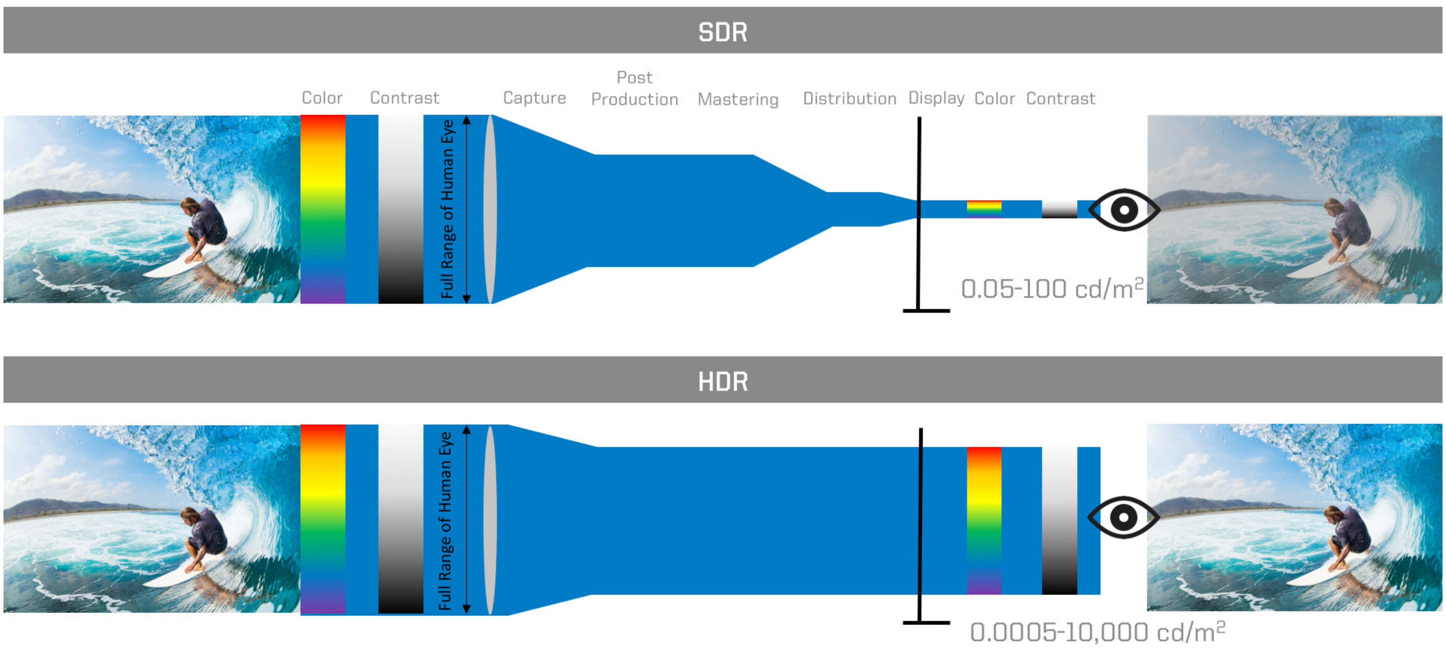

The diagram above shows a representation of the full range of light the human eye is capable of seeing. HDR is able to retain a greater range of light than SDR from capture to final distribution. Pictures are more vibrant and life-like to the viewer, regardless of screen size.

The luminance encoded in SDR signals is relative – at 100 percent, it basically tells the display to show its whitest white. In contrast, HDR codes the absolute value of the luminance, using a non-linear transfer function based on what the eye can perceive. This is the SMPTE 2084 Perceptual Quantizer (PQ) transfer.

Beyond luminance, HDR also features wide color gamut (WCG), which has more color information than standard HD signals. The set of colors a signal uses or a display can show is called the "color space." There are three defined color spaces in use today: ITU-R BT.709, which is considered the standard color space for HD; DCI-P3, which is the standard for digital cinema; and ITU-R BT.2020, which is the UHDTV standard.

No commercial monitor today can display the full Rec. 2020 color space, but a UHDTV with HDR and WCG will show more than the entire Rec. 709 color space and at least 90 percent of the DCI-P3 color space. But what happens when a display is fed an HDR signal it cannot display because the luminance and/or color

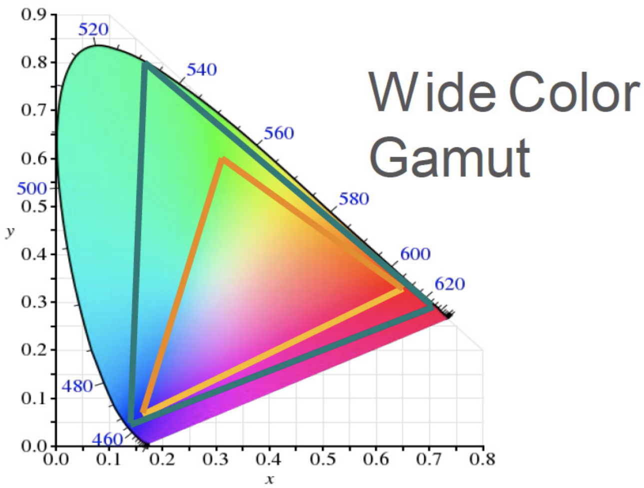

This graph is a representation of the color gamut. The orange triangle represents REC.709 or 35.9% coverage. The green triangle represents REC.2020 which is 75.8% coverage.

are out of range? The monitor must create an image as close as possible to the original source material – and to help the monitor do this job, metadata may be included in the stream. Metadata helps the monitor adapt the absolute luminance of the HDR signal to its capabilities.

Static Vs. Dynamic Metadata

Metadata can be static or dynamic. Static metadata, as the name implies, is fixed for the duration of the content, and provides only "general" information. The basis for static metadata is SMPTE 2086. When content is created, it is mastered on a reference display by the director/colorist. The static metadata describes the characteristics of this mastering display, so that the monitor currently playing it can best approximate the mastering display based on its capabilities. SMPTE 2086 static metadata includes parameters such as the color and luminance range of the mastering display, the color space used, and the transfer characteristics. SMPTE 2086 was augmented by CTA to include additional parameters such as the Maximum Content Light Level (so the monitor will know the "brightest" pixel in the content) and the Maximum Frame-Average Light Level.

The objective of dynamic metadata is the same, but it changes from frame to frame. Static metadata can be seen as sort of an average over the content, while dynamic metadata is tailored to each individual frame. This dynamic metadata is one of the ways the various HDR standards differ.

HDR Standards

There are a number of competing HDR standards available today, and most of them provide some support for simultaneous SDR and HDR support. From a high-level point of view, they can be classified as static (using static metadata or no metadata) and dynamic (using dynamic metadata).

The baseline HDR support starts with the SMPTE 2084 PQ transfer function, using 10-bit or 12-bit samples. This is the basis for most HDR standards, and by itself does not include any support for SDR. The HDR10 standard is simply the combination of SMPTE 2084 with 10-bit samples and SMPTE 2086 static metadata; this combination is also standardized in ATSC A/341.

Static vs Dynamic Metadata

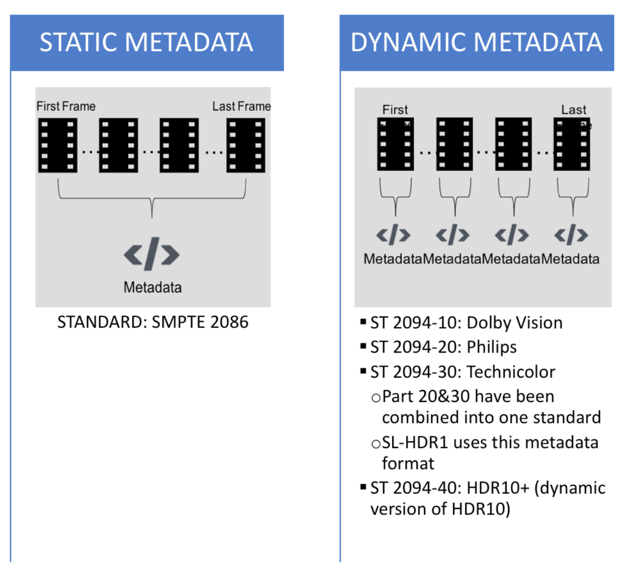

Static metadata provides only one metadata value while dynamic metadata provides a different metadata value for each frame of video.

One interesting static HDR standard that is in wide use today is Hybrid Log-Gamma (HLG), which is not based on SMPTE 2084 PQ. It is an attempt to use a backward-compatible transfer curve that will "work" with both SDR and HDR monitors without any metadata. At the low luminance levels, it matches SDR, so an HLG signal applied to an SDR monitor will "look OK", which an HDR monitor will show the improved ranges at the higher luminance levels. HLG trades off simplicity (same signal everywhere, no metadata processing) with quality (it is not as good as the dynamic metadata options). HLG is standardized in ARIB STD B-67, ITU-R BT.2100, and ATSC 3.0.

The dynamic HDR standards all start from a PQ base layer, with metadata defined in ST 2094-1. The most relevant ones are SMPTE 2094-10 (Dolby Vision), SMPTE 2094-40 (HDR10+, a dynamic version of HDR10), and SL-HDR1.

HDR-to-SDR Conversion Options

The basic operation of various dynamic standards can be understood as transmitting "an image plus instructions" that can be processed by a monitor. What varies is what you start from. In Dolby Vision and HDR10+, you start with an HDR image, and the "instructions" allow the mapping of that HDR image to any monitor, all the way down to an SDR monitor. While it is possible for an end device to generate SDR from this HDR signal, that end device needs to understand and process the metadata in order to do so.

With SL-HDR1, the opposite happens. What is transmitted is a standard SDR signal, and the metadata allows a compatible device to reconstruct the original HDR signal (or any intermediate level suitable for its capabilities). This is the ideal way to support legacy SDR devices – they will just ignore the metadata (because they do not support it) and simply display the SDR image. This is conceptually the same as what was done when analog color TV was introduced. The signal was the standard black-and-white content, with the color information added "on the side". A black-and-white TV would understand the signal and show the black-and-white version, while a newer color TV would extract the color information and show a color picture.
To support a mix of SDR and HDR devices, broadcasters have several options. First is simply simulcasting two or more versions of the same signal, one in HDR and one in SDR. Simulcast is well suited for OTT distribution, where multiple versions of the same content are already required. It is also the only real solution for supporting legacy 8-bit devices, such as AVC 4:2:0 decoders, because HDR requires 10-bit performance to avoid posterization.
HLG is another option, as the same signal is compatible with both HDR and SDR. However, this can negatively impact the quality of the

SL-HDR1 Workflow

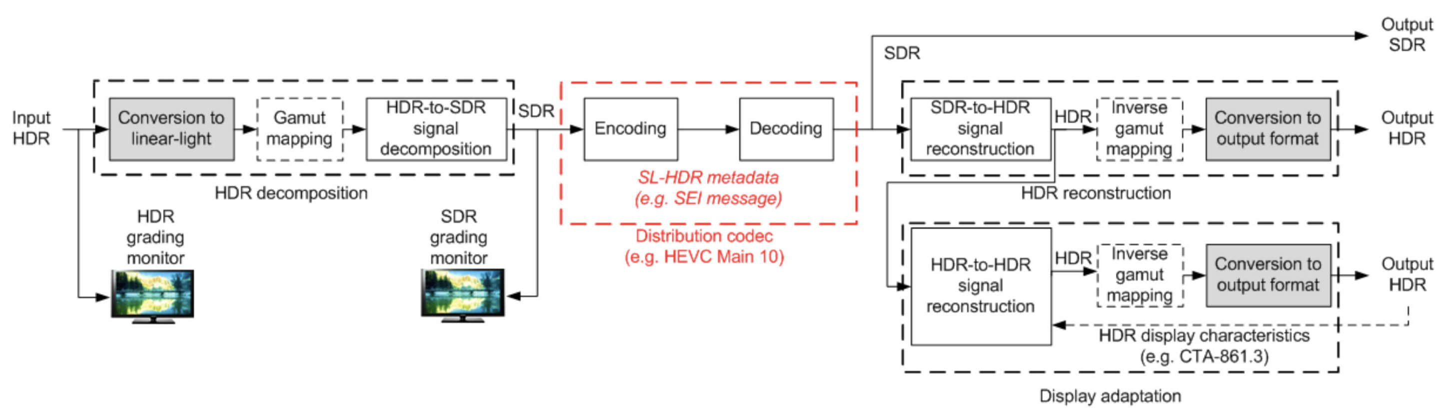

Source: ETSI TS 103 433-1
The SL-HDR1 workflow allows for simultaneous HDR and SDR delivery with just a single SDR output with metadata from the encoder. A downstream receiver determines HDR or SDR playout.

HDR signal, unless the distributor controls the production. Other dynamic modes transmit HDR with metadata that can be used to reconstruct SDR, but it requires a receiving device that understands HDR and metadata.

A better choice is SL-HDR1, which transmits a very good quality SDR signal with metadata to reconstruct the HDR image – and automatically delivers a full HDR experience to compatible monitors and/or set-top boxes. SL-HDR1 is defined in ETSI TS 103 433-1, and was approved in early 2018 as an amendment to ATSC A/341. Metadata is carried through inside the video elementary streams as SEI messages in compressed streams (H.264/H.265), or carried in the VANC using SMPTE 2108 for SDI.

For scalable HEVC (SHVC), the A/341 standard calls for two spatial layers – base and enhancement – and the SL-HDR1 metadata may be included in either layer. The spatial resolution of the enhancement layer is up to three times that of the base layer. If the SL-HDR1 metadata is in the base layer, it applies to both layers, but if it is present in the enhancement layer, it applies only to it. This gives the flexibility to the broadcaster to have different signal levels. For example, they can broadcast a free SDR base layer at lower resolution, and add a premium HDR layer at higher resolution.

Case Study: Spectrum Networks

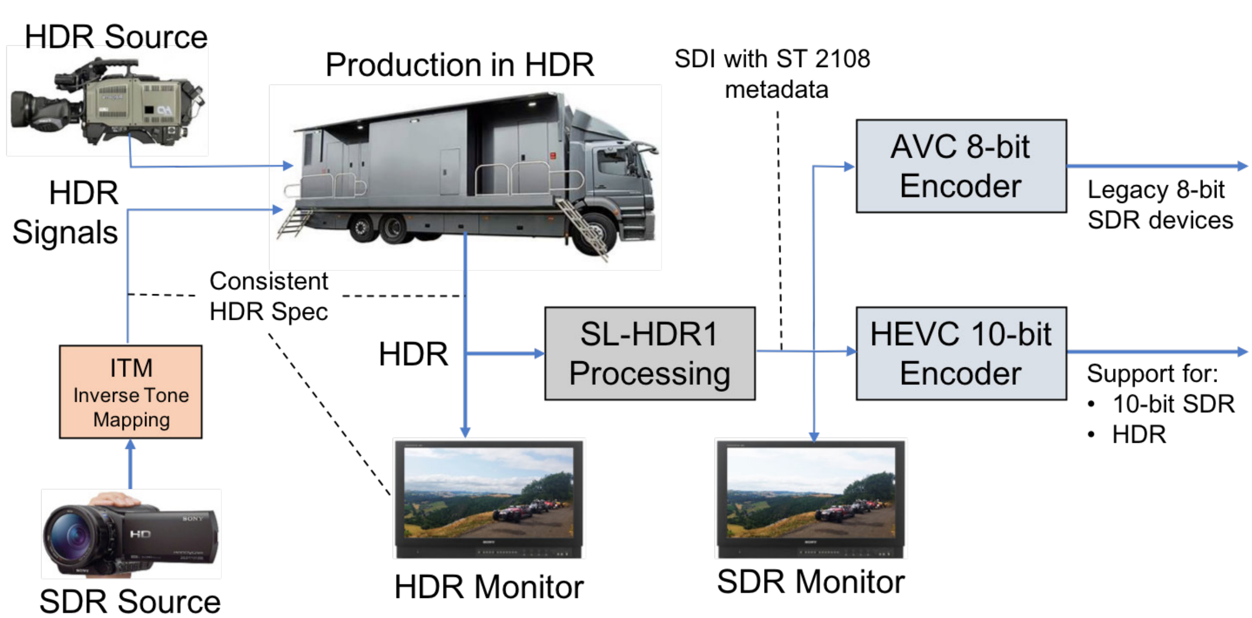

Spectrum Networks was able to ingest both HDR and SDR feeds with a single cost-effective truck, while providing viewers with both high quality HDR and SDR content.

The viability of SL-HDR1 conversion was proven in Summer 2017, when Spectrum Networks produced and broadcast a Los Angeles Dodgers home baseball game in HDR for MLB Network. All SDR sources were upconverted to HDR. An SDR downconversion was used to check the SDR feed. The production generated HDR and SDR versions simultaneously using a single production environment, distributing SDR, SL-HDR1, and HDR10 (converted from SL-HDR1) versions to audiences.

Since then, SL-HDR1 processing continues to be used to produce live sports production. It was used successfully by AT&T to deliver 4K HDR imagery of its live coverage of “MLB Network Showcase” games. NBC Sports Group also delivered all seven home games of the Notre Dame Fighting Irish football team’s 2018 season in 4K HDR for AT&T’s DirecTV subscribers. The football games were produced on site in 1080p, and upconverted live to 4K HDR.

The Cobalt Solutions

Central to these productions is a workflow built around Cobalt Digital's flagship 9904-UDX-4K up/down/cross converter and image processor for openGear® frames. Rather than hire separate trucks for HDR and SDR productions of the same event, the 9904 can be used to convert SDR material to HDR (using a process called "Inverse Tone Mapping") in real-time. This way, the whole production chain can be in HDR. The processing card uses the Technicolor HDR Intelligent Tone Management (ITM) software to provide this mapping function. The Technicolor technology defines how the available dynamic range is used, managing brighter lights and darker shadows to deliver sharper, more realistic images.

For HDR output formats, the 9904 supports HLG, Sony S-LOG3, and PQ10, with 4K resolution support via 12G or quad 3G inputs. HDMI 2.0 output is also available with HDR InfoFrame support for monitoring. It can also apply an arbitrary static 3D LUT to a signal for format conversion, color space conversion, or static ITM. The 9904 is compatible with Pomfort LiveGrade and Wowow WonderLook on-set look management software.

Other Cobalt card-based products that can be part of the SL-HDR1 workflow include the 9223 3G/HD/SD MPEG-4 encoder series, which supports legacy 8-bit, 4:2:0 SDR devices. The 9992-ENC encoder supports H.264/H.265 8-bit and 10-bit applications, with resolutions up to 4Kp60.

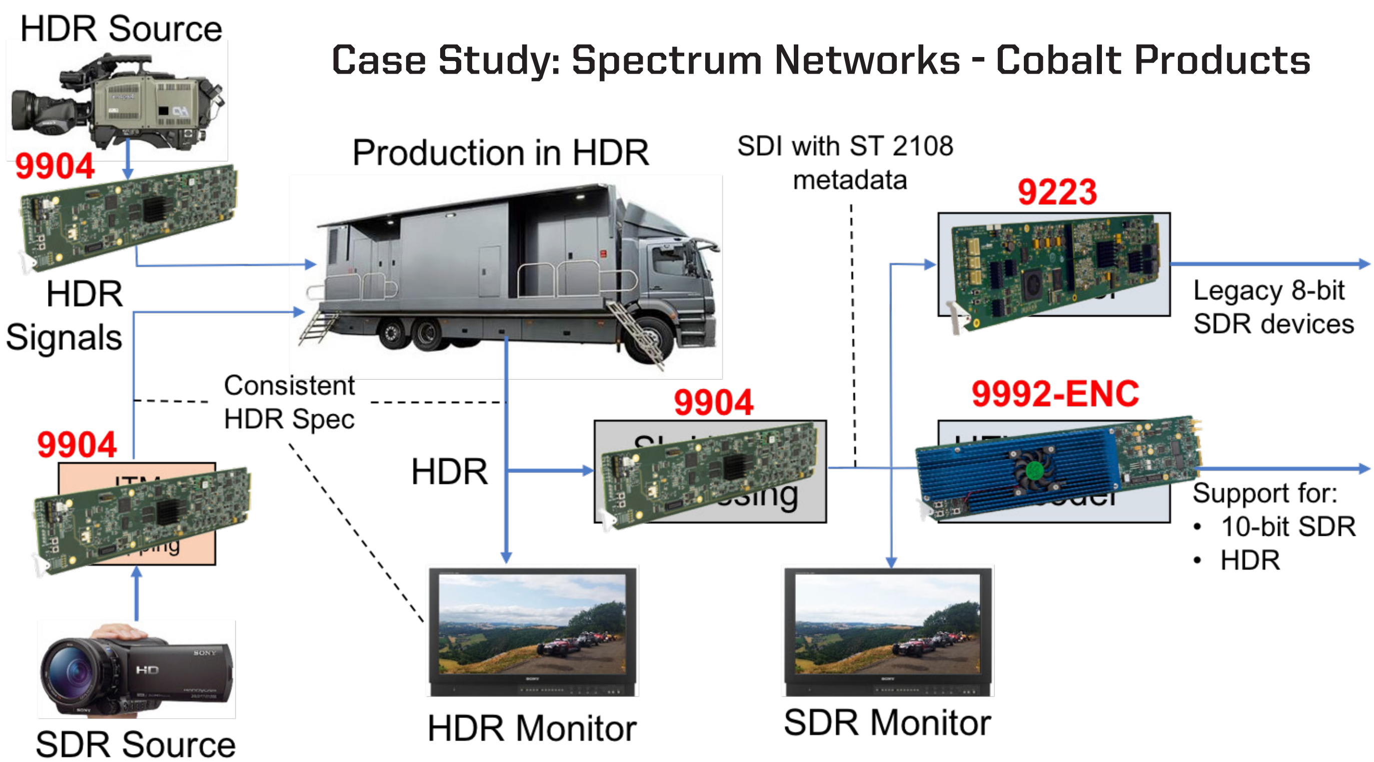

Cobalt Digital Inc. designs and manufactures award-winning edge devices that help live video production and master control clients transition to IP, 4K, HDR, the cloud, and beyond. As a founding partner in the openGear® initiative and proud member of SMPTE, Cobalt also offers a best-of-breed interoperability platform that simplifies technological adoption. Distributed through a worldwide network of dealers, system integrators, and other partners, Cobalt Digital products are backed with a five-year warranty.

WWW.COBALTDIGITAL.COM