The Brave New World of Loudness Control (Part 1)

New Post-Production Workflows Transcend Compliance for Enhanced Audio Creativity

By
Jon Schorah
(NUGEN Audio)

Coalescing around the ITU-R BS.1770 standard—Algorithms to Measure Audio Program Loudness and True-Peak Audio Level—loudness regulations or recommendations are in place across the US as well as South America, Europe, Japan, and much of the world. No matter where you produce audio and for whom, it's almost certain the broadcast will need to be loudness-compliant. It's just a question of how to make it happen.

As loudness management becomes the required norm, the responsibility for compliance increasingly falls on audio post-production engineers. But there's some good news. Since several world regions have been delivering loudness-compliant broadcasts for two years or more, a methodology for best practice is beginning to emerge. The solutions for loudness control are not only maturing, but also becoming versatile tools in the engineer's toolbox.

Initially, loudness compliance may appear to be just another task added to the seemingly ever-growing list of tickboxes engineers must check before delivery. But if loudness normalization is correctly integrated, it can become more than a means of avoiding consumer complaints and potential fines. It can reintroduce creative freedoms lost in the old peak normalization paradigm and become a tool to improve broadcast audio quality.

Loudness and Post-Production: A Creative Marriage

At first glance, playout processing is an obvious solution to loudness compliance. By definition, playout processing achieves loudness normalization as an afterthought by adjusting station output to ensure compliance after post-production. However,

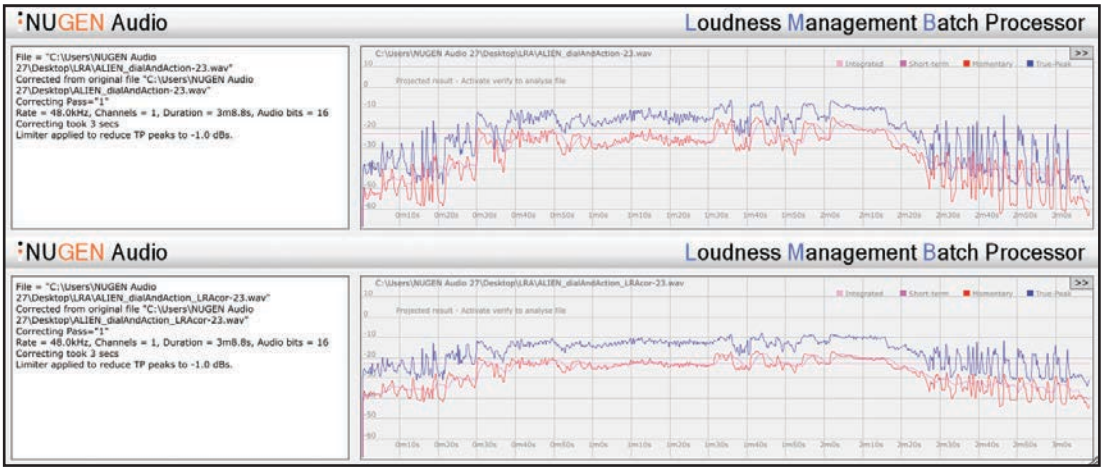

Figure 1: This section of a feature film audio has been repurposed for TV by reducing the dynamic range. The top image shows the original signal analysis.

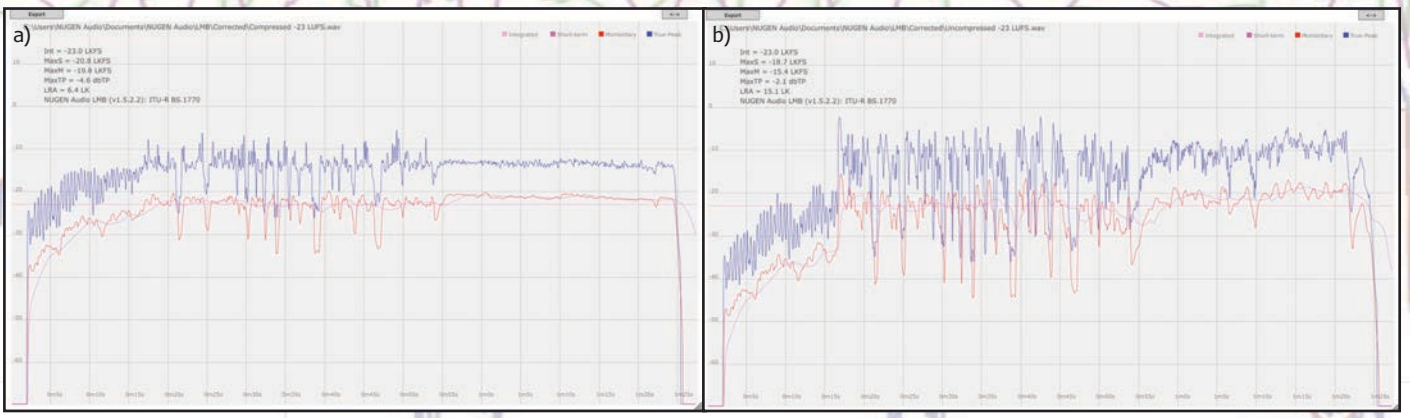

Figure 2: A section of heavily compressed audio has been normalized to –23 LUFS (a). The same section of audio is shown with light compression normalized to –23 LUFS (b). A comparison clearly shows that instead of producing a "louder" result, much of the heavily compressed signal is below that of the corresponding lightly compressed alternative after loudness normalization.

experience has shown that playout processors are not particularly good at solving loudness issues in a consumer-satisfying manner. They can even introduce their own new loudness problems. In regions where loudness recommendations have been in place for more than a year, there is a clear movement away from correction after the fact.

A better solution would be to consider loudness compliance as part of the creative process during pre-production and even at acquisition. In this scenario, simply pushing the mix against the limiter is no longer a viable option. Instead, incorporate loudness compliance into an experienced post-production professional's creative decision-making process. In general, noncompliant material that falls outside a tolerance margin is rejected. Near-compliant audio can be satisfactorily corrected with loudness batch processing that brings it into compliance with a gain offset and possibly some true-peak limiting. Here, the playout processor's role shifts to that of an error-handling stop-gap and is bypassed with the delivery of compliant material.

Instead of adding another complicating factor to the mixing process, the right loudness tools provide an engineer more creative freedom. Compressing a mix to achieve a consistently loud level under loudness normalization will cause the audio to be turned down. So the best way to achieve a mix that stands out in the crowd is to creatively engage with the content and make the most of the available dynamics.

Loudness normalization coupled with the use of true-peak maximum levels enables you to creatively use dynamic range and increase headroom. When experienced operators ensure compliance in post-production, the playout processor becomes largely inactive. These are great developments, especially for professionals involved in short-form production.

Loudness can be considered in pre-production and acquisition. Normalizing archive and library material and ensuring field recordings and outside broadcasting (OB) sources are already compliant speeds the production process, enabling a faster turnaround while ensuring the audio remains faithful to its original context. This can be especially

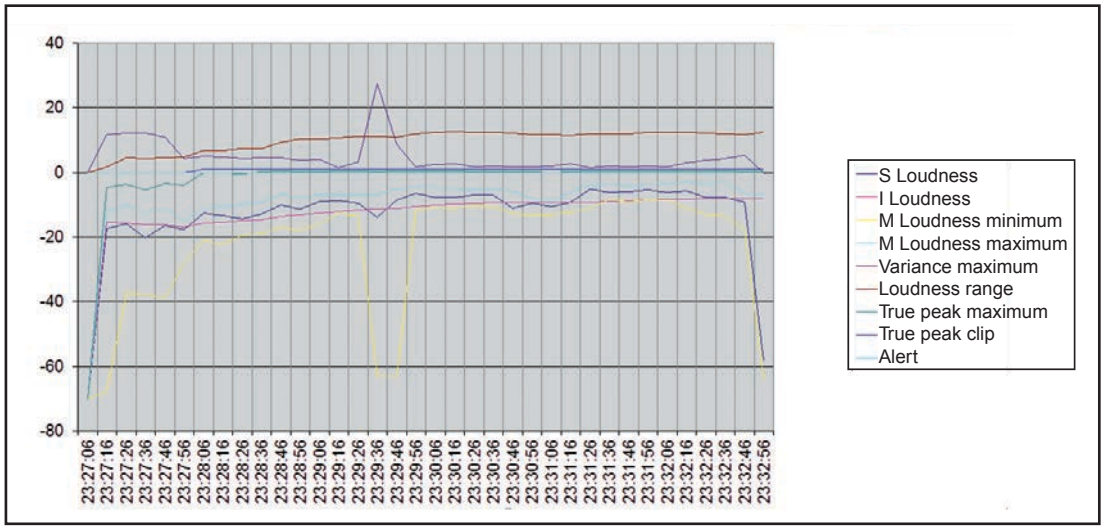

Figure 3: The typical log file output from a real-time loudness meter (VisLM-H) shows the variation of loudness parameters over time.

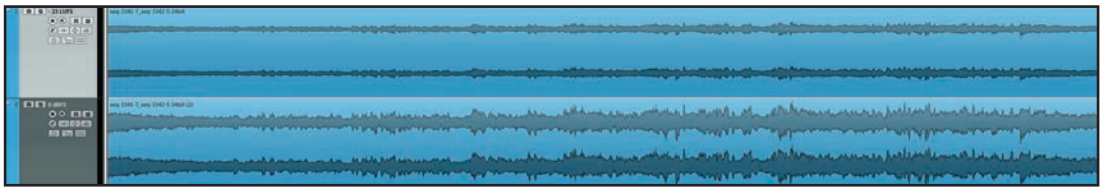

Figure 4: The same section of audio is normalized to –23 LUFS and to 0 dBFS to show loudness vs. peak normalization.

important during live sports events, breaking news broadcasts, event coverage, and studio interviews.

Building a Post-Production Loudness Workflow

New workflows call for new tools. Fortunately, companies are developing products designed specifically for post-production engineers. Intuitive audio-editing tools (e.g., real-time metering, offline correction, and loudness-compliant limiting) enable post-production editors to put their creative expertise to work while ensuring compliance.

Ears are the best tools when it comes to making creative decisions in audio post-production. That rule also generally applies to loudness normalization. The new loudness standards hold the potential for increased dynamic range and contrast. But computers also play a key role with their ability to take measurements and make smaller adjustments to get things exactly on target. Computers can work quickly, saving considerable time near the end of the process. With these tools in mind, a post-production workflow begins to emerge.

Clear, intuitive loudness metering is the key to delivering high-quality, loudness-compliant audio. Because the new loudness measurements are designed to correspond to the human ear, a good engineer can almost mix sound without a meter. It's possible for skilled engineers in a calibrated room to simply occasionally glance at the meter during the creative process to maintain their bearings or to check something in particular. But in the end, even experts need to confirm that they've met target values. Likewise, anyone looking to push creative boundaries needs checks to ensure their work is compliant. With visual meters, editors can keep an eye on the meter and loudness profile while relying on their trained ears to make most of their decisions.

Another important factor in this workflow is a high-quality true-peak limiter that can handle the new standard's intersample true-peak requirement, which is something traditional sample-peak limiters cannot do. It's tempting for engineers to use their existing sample peak limiters with a setting that would yield results "safe enough" to be compliant with the loudness standards' true-peak measure. But, those who follow this practice do so at their peril. Simply put, it's impossible to arrive at an accurate true-peak reading with a sample peak limiter because the measurements are different. What may seem like safe settings on a sample peak limiter would not guarantee compliance.

Therefore, the best true-peak limiters offer a true brick-wall solution, measuring inter-sample peaks and enabling the user to define the audio output's true-peak limit (rather than the more traditional threshold control at which limiting begins to take effect). Based on ITU-R BS.1770's standardized true-peak algorithms, these tools are suitable for controlling audio for post-production and broadcast applications. True-peak limiting can also be used to ensure that downstream codecs (e.g., MP3, AAC, and others) do not introduce distortion into the signal.

Once a mix is more or less loudness compliant,

Figure 5: Production to loudness standards are shown within a typical nonlinear editor (NLE).

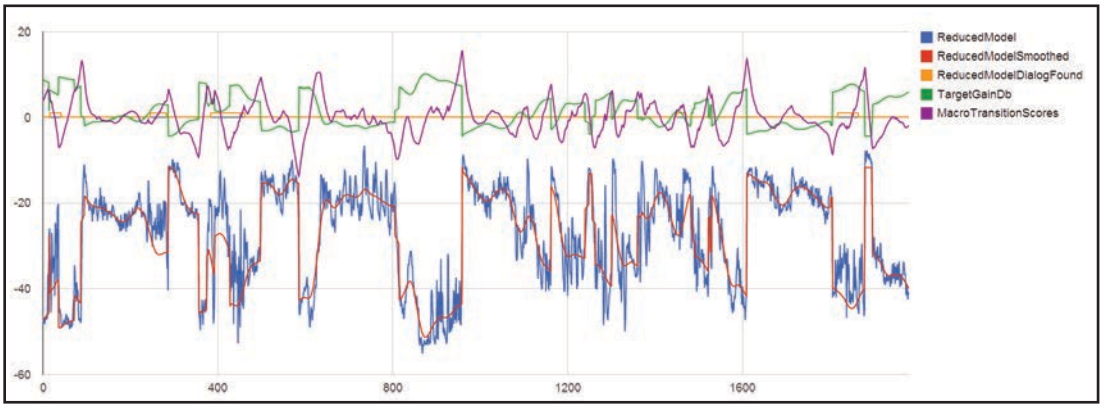

Figure 6: There are complexities involved in reducing the dynamic range of a feature film while continuing to respect dialogue clarity.

editors can use offline tools to fine tune the mix and speed up the last part of the normalization process. These timesaving tools can be plugged into the editing environment to bring a mix into line quickly, correcting any true-peak overshoots along the way.

Batch analysis is another highly useful tool for busy post-production operations, enabling them to automate part of their loudness processing. Acting as a rapid fail-safe system and internal QA component, a batch processor can automatically assess files for compliance and correct or reject as needed.

An Expanded Role for Loudness

Loudness measurement doesn't have to end when broadcast criteria are met. In regions where loudness compliance is an accepted part of the audio workflow, the same normalization tools can also be employed in several areas that go beyond technical compliance to support new and improved production techniques.

As I mentioned, loudness consideration during acquisition is a significant timesaving pre-production technique that brings audio into the editing suite at the right time. Another area in which loudness can play an important role is dialog clarity. Even today, mixes are occasionally broadcast with the background music too loud, which makes dialog indistinct and results in viewer complaints. Using a meter to preserve loudness separation for dialog above other mix components can help guard against these mistakes. Measuring the loudness of background music beds and FX spots can also help maintain consistency from section to section. In sessions that require significant complex editing, loudness normalization can quickly match dialog levels to a far more useful control than 0 dBFS (i.e., the maximum possible digital level).

Audio libraries can also benefit from loudness normalization, ensuring that audio is always internally consistent and available at an expected level. Consideration of the loudness range (LRA) parameter can also be useful when mixing material for a specific target device.

Figure 7: A sample-by-sample analysis shows a True-peak over.

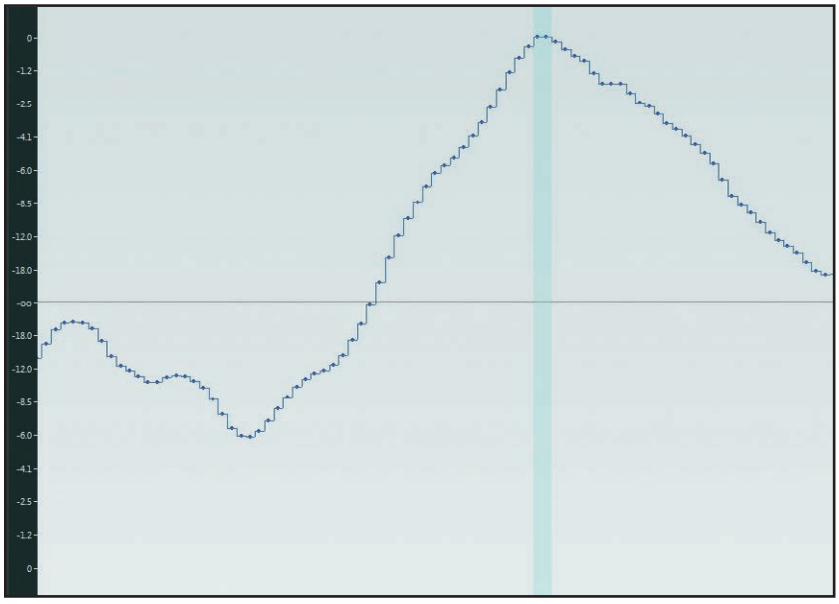

Work Still To Be Done

Fortunately, the various international loudness recommendations are based on the same International Telecommunication Union (ITU) standard, so there is general agreement within the industry about how to approach loudness control. Even so, some issues still need resolving.

One problem that can arise in a maturing market relates to internal loudness jumps, which can cause viewer irritation or discomfort. In the pressurized advertising environment, commercial imperatives demand every method of capturing viewer attention is employed. Therefore, a clever mix engineer may carefully mix a spot with a long quiet section, which would enable a very loud burst of audio without affecting the spot's overall loudness compliance. How to handle this new consumer annoyance is far from standardized. Some regions have introduced additional loudness constraints for commercials that

go beyond the integrated program loudness (e.g., maximum momentary or short-term loudness) to avoid this exploit, but they're walking a fine line between loudness control and over-specification that can trample on desirable creative freedoms.

A more technical issue is the difference between a 5.1 mix and its corresponding downmix. It is common for the downmix to differ slightly in loudness from the 5.1 mix. However, the difference can be in either direction. This can further confuse the issue and preclude a simple offset as a viable solution. Similar situations arise with dual-language, multi-mono stereos, in which a consumer's television can produce an unexpected 3-dB loudness jump depending on the configuration. Relying on the metadata is one solution, but this only works if the metadata is accurate and the appropriate device is capable of properly reading and responding.

Work also needs to be done in the application of loudness compliance to other audio fields. The loudness standards for broadcast were not designed with radio, game audio, music production, or film in mind, and there is room for further research into how these areas may benefit and what specific requirements may be needed.

For example, loudness measures are only defined for up to 5.1 audio, but much film and game sound is now produced in 7.1, which is not covered by the ITU-R BS.1770 definition. Now that iTunes Radio has incorporated a sound-check algorithm for level harmonization, loudness considerations during music production are also coming to the forefront.

What's Next?

As a product designer, this is one of the most common questions I get in a mature loudness market. The answer lies in two areas since new developments are not just about the tools. First, we will see new applications arising from the continuing education process and a deepening understanding of existing tools. This will lead to a demand for enhancements to existing tools (e.g., warnings when loudness differentials between differing signals or signal types become too large or too close). Multiple-stream loudness metering could be employed in a broadcast studio to monitor the microphone levels during a live-panel interview. If the meters are aware of a preferred target range, they could alert the operator to potential dialog clarity issues if signals were to become too divergent.

Similarly, the opposite could be applied in an OB production if the crowd noise at a sporting event rises too close to the commentary level. At the moment, most automated alerts are based around compliance issues, but as products develop we will see an increased use of flexible alerts to identify potential clarity concerns based on a local set of preferences. With awareness of the overall program loudness target, the producer could retain clarity and meet loudness targets in real-time situations that require fast and accurate decision making.

Looking further ahead, new loudness algorithms that measure the loudness of dialog in audio relative to the overall program loudness could be developed to enhance the repurposing of film audio for television. Currently, traditional limiting takes no account of the level of the dialog and tends to push the dialog too low in the mix as the dynamic range is reduced for reasonable television broadcast. An algorithm that is aware of the dialog level would be able to compress around this anchor and preserve the dialog level within the resulting repurposed work.

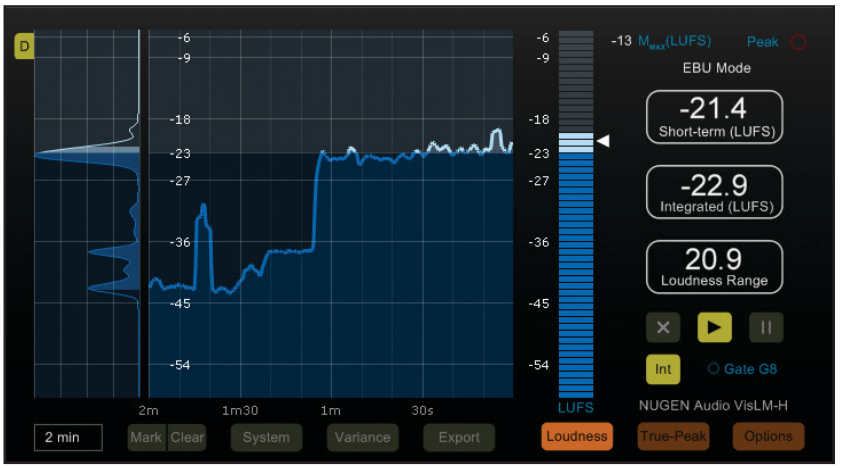

Figure 8: The NUGEN Audio VisLM-H provides a simple standard-compliant way to measure, compare, and contrast loudness during production, broadcast, and post production.

Workflow Changes Are Necessary

Establishing a post-production workflow that integrates loudness normalization leads to improved audio quality, with compliance becoming an integrated and natural part of the creative process. The practice of loudness compliance leads to better audio consistency and enables a greater dynamic range with the introduction of more headroom resulting in the opportunity for more creative expression.

Thanks to loudness normalization, there are fewer complaints from consumers, which means that the changes are headed in the right direction so far. Now it's time to start refining solutions for loudness control and build on the solid foundation already established.

Ultimately and ideally, loudness will become a primary consideration during production. As tools improve, loudness parameters and transferable objective measures can be used to check whether audio is compliant and target-appropriate. At the same time, audio engineers can use these tools to produce better-sounding, more creative mixes.

About the Author

Jon Schorah is the creative director and co-founder of NUGEN Audio, one of the world's leading manufacturers of loudness products. Jon has a background in mastering and engineering and has considerable experience in wider aspects of the industry. A 1992 Leeds University (UK) graduate, in recent years Jon has focused on product design with a particular interest in the usability and workflow aspects of audio software.