Loudness Measurements and the CALM Act

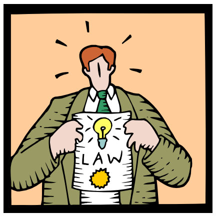

The recent passage of the CALM act has made loudness management a major concern for broadcasters and those that supply them with programming and advertising. Most broadcasters and content creators understand that compliance will come at a cost, in capital expenditures, changes in their workflow and additional record keeping in the form of logs and reports. However, before investing in any equipment or process it is essential to understand the legislation and its likely implications. This note explains the CALM act and its ramifications, and suggests cost effective solutions to conformance issues.

Background

The CALM act is a response to the increasing mismatch between the loudness of programs and commercials. In the days of analog television this mismatch was mitigated by the use of dynamic range compressors. These were commonly employed at the transmitter input and dynamically adjusted gain to insure that all material was broadcast at approximately the same level.

Though good compressors worked reasonably well, their effectiveness could be compromised by advertisers who manipulated their content to obtain a subjectively louder result. This was possible in part because the compressors didn't really sense audio loudness. It was also due to the inherent tradeoff in any compressor between maintaining constant output and retaining some semblance of program dynamics. Since compressors had to allow some program dynamics through to sound decent there was the opportunity for advertisers to manipulate their contents dynamics to sound louder than normal program material.

The situation was annoying to most listeners but remained tolerable until the transition to digital television sound. Several factors contributed to the drastic worsening after the transition.

- Most of the compressors in use were analog devices and only handled stereo signals. Digital TV audio requires 5.1 channel equipment which interfaces digitally. The old gear was retired and equivalent 5.1 channel gear with digital I/O was not generally available.
- The wider dynamic range available in the ATSC system made many broadcasters think that dynamic range wasn't an issue and compressors weren't necessary.
- The ATSC system included a new parameter called dialnorm which was easily misunderstood, ignored by others and consciously misused by savvy advertisers.

Dialnorm

Dialnorm is metadata which directly controls the playback level over a 31 dB range. The name is a misnomer since it does not adjust only the dialog level heard by the viewer but instead sets the level of all channels and content during reproduction. The dialnorm value was intended to represent the loudness of dialog in the content. Because the viewer's decoder uses this to control its gain the overall program loudness and the dialog loudness remain constant regardless of the content.

Loudness Measurement

Audio level has been characterized using many different techniques, VU meters, various standard PPMs, and even oscilloscopes. The transition to DTV and the ensuing doubling of available dynamic range of audio content can result in a huge difference in the perceived loudness of that very content. Prior to DTV the available techniques for measuring level left little need to accurately assess loudness due to its restricted range. Enter the attempt to mitigate this situation by developing a new measurement paradigm and standardizing it across the entire spectrum from content creation through transmission, Loudness measurement (as standardized in ITU BS.1770) requires modeling the non-flat frequency response of the ear and measuring the resulting signal energy. The correlation between loudness measurements and subjective loudness is not perfect, but is generally good enough to make them practical and useful. A detailed explanation of BS.1770 loudness measurement and a comparison to conventional audio level measurements can be found in the Qualis Audio tech note "Loudness Measurement".

CALM Act requirements

The CALM Act mandates the FCC to write rules enforcing "the ''Recommended Practice: Techniques for Establishing and Maintaining Audio Loudness for Digital Television'' (A/85), *and any successor thereto*, approved by the Advanced Television Systems Committee, only insofar as such recommended practice concerns the transmission of commercial advertisements by a television broadcast station, cable operator, or other multichannel video programming distributor.".

This wording makes the law apply to virtually everyone that delivers commercial content to a viewer that doesn't travel on a physical media such as DVD. It doesn't matter if it came from someone else and you only pass it along, if you transmit it you are liable. Whether you can extract reimbursement for FCC fines from entities further upstream is your challenge.

Note that the "*and any successor thereto*" phrase effectively makes the ATSC the regulatory authority and the FCC the enforcer. The ATSC recommended practice sets the expected behavior of broadcasters. The enforcement and fine assessment procedures will be established by the FCC. If the ATSC changes A/85, the requirements the FCC puts on broadcasters automatically change. Presumably the specifics of enforcement will remain unchanged.

The enforcement procedures for the CALM Act won't be known until the FCC issues it regulations. However, it's a good bet they will follow the approach used for closed caption regulations, obscenity and other requirements. The FCC will not police the airwaves. Instead, they wait for a viewer complaint. When a complaint is lodged the FCC opens an investigation. Unlike a classic police investigation, they simply notify the broadcaster of the alleged infraction, its date and time. It is up to the broadcaster to provide evidence to counter the complaint. Normally this evidence will be a log of the broadcast at issue. If the log proves compliance with FCC regulations the complaint is dismissed. If the log does not prove compliance or if a log is not available the broadcaster is considered guilty and a fine is assessed.

In this scenario logging becomes a bottom line issue. Prove your broadcast complied with ATSC A/85 and you are OK. If you can't, you are fined. The cost of one fine is likely to exceed that of a loudness monitoring/logging solution, easily justifying the investment in one.

ATSC A/85 requirements

The ATSC A/85 document is a recommended practice, not a standard per se. However, its inclusion in legislation effectively elevates it to a standard. As such it is important to understand what it recommends.

- Loudness of all segments must be measured according to the technique specified in ITU BS.1770.
- The dialnorm value broadcast with each segment must match the measured loudness within +/- 2LU.
- It is recommended that dialnorm be set to -24 but any value may be used as long as it matches the segment loudness.

Operating with a variable dialnorm value which tracks the segment loudness requires measuring each segment and storing the result in file metadata. This isn't possible for live broadcasts since the program is not available in advance and hence requires the use of a fixed dialnorm.

An insidious problem occurs when something goes wrong in the path from file metadata to the broadcast metadata. The dialnorm field will contain something. If it isn't the file metadata value it has a 1 in 31 chance of being correct. When a mistake creates exposure to a fine, minimizing mistakes becomes the focus.

The likely outcome of this situation is that broadcasters will view variable dialnorm as impractical and dangerous to implement. They will operate with the recommended fixed dialnorm of -24 and normalize all content to a loudness of -24. When live material is broadcast the mix engineer will target a -24 value for loudness. Using this approach there is no requirement to pass metadata from the file to the encoder. Consequently there is no risk of this path being broken.

Real Time Monitoring and Logging

The audio portion of a typical television broadcast consists of sequential programs which are separated and/or interrupted by commercial advertisements. The standards specify, and commercial equipment implements, the ability to start, stop, pause and reset loudness measurements. This allows the user to begin measurement of a selected program, pause it during measurements of commercial breaks, resume measurement when the program resumes, stop the measurement when the program concludes and reset the meter to allow another measurement. This control can be effected manually or through hardware connections to the measurement equipment. This suits measurement of file based content but has limitations when applied to real-time measurement.

Real-time monitoring and logging requires independent assessment of program material and of the commercials which separate or interrupt it. Performing independent assessment of both the program and the commercials requires multiple measurement meters or convoluted and error prone manipulation of the data provided by a single meter.

The problem posed by a typical broadcast audio stream is illustrated in Figure 1. The signal consists of multiple programs broadcast sequentially, identified in the figure as P1, P2, etc. These programs are interrupted by commercials, identified in the figure as C1, C2, etc.

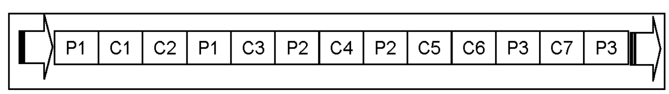

Figure 1 A typical broadcast audio stream

The interruptions may consist of a single commercial or several sequential commercials as illustrated in the figure. When one program concludes it is typically followed by one or more commercials, identified in the figure in the same manner as the commercials cited earlier. When the inter-program commercial break concludes the next program is transmitted. It is interrupted and followed by commercials in the same manner as the preceding program.

For a loudness meter to measure and log this stream it is generally told when to begin a measurement by a start signal provided either manually or automatically by the user. When a commercial break occurs the meter is paused using a second control signal provided by the user. At the program conclusion the measurement is stopped and the measurement value is obtained. If the intervening commercials are to be measured and logged, an additional meter is required with its own set of start and stop controls and a means of merging the log data with the other meter.

The audio stream may be viewed as consisting of a program stream and a commercial stream. One or the other is always being broadcast. There are some relationships between these two streams that allow simplification of the measurements. These are:

1. The programs are always sequential.
2. Programs may be interrupted by commercials.
3. The commercials are always sequential.
4. Commercials are not interrupted.
5. Commercials may occur between programs.

Since commercials are not interrupted, it is not necessary to pause measurement of the commercials. If a pause occurs during a program it must be for a commercial. If the meter understands this we no longer need a pause control, just a control to select “program” or “commercial”. Since there is always content being broadcast the start/stop control can be replaced with a “measurement reset” or “new” control.

Qualis Audio Implementation

Figure 2 illustrates how this is implemented. The measurement is controlled by two signals; a stream select and a measure/reset. These can be logic level inputs from a playout server, dedicated buttons or soft buttons on the instruments GUI. Similar control commands received over the instruments Ethernet port are possible.

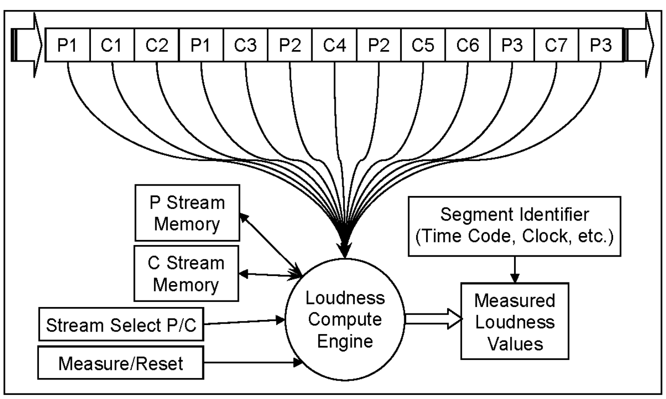

Figure 2 The Qualis Audio dual loudness measurement architecture

The stream select signal depicted in figure 2 causes the loudness measurement engine to measure the program stream (Pn) when in one state, for example when high, and the commercial stream (Cn) when in the other state, in this example when low. The engine always measures one stream or the other. However, there is additional functionality built into the way the measurement engine uses this control signal. When the signal transitions from the P state to the C state the program stream measurement is paused. The program stream measurement memory is not cleared or affected in any way. When the signal transitions from the C state to the P state the current loudness measurement is saved and logged as a final result and written to the results memory. The commercial stream memory is then cleared to prepare for another measurement and the loudness measurement engine now continues measurement of the ongoing program stream.

The measure/reset input causes the loudness measurement engine to save and log the current measurement value as a final result and clear the stream memory to begin a new measurement. This is done regardless of which stream is being measured. If the engine is measuring the program stream when reset is asserted the current reading is written to the result memory and the program stream memory is cleared. If it is measuring the commercial stream when reset is asserted the current reading is written to the result memory and the commercial stream memory is cleared.

The loudness measurement engine also keeps a running Short Term measurement which is not affected by the control lines. This continuous measurement tracks listener perception of the stream in real time and provides assessment of trends in the loudness of the current stream.

These readings are displayed on a dual bar-graph loudness display giving a complete loudness picture.

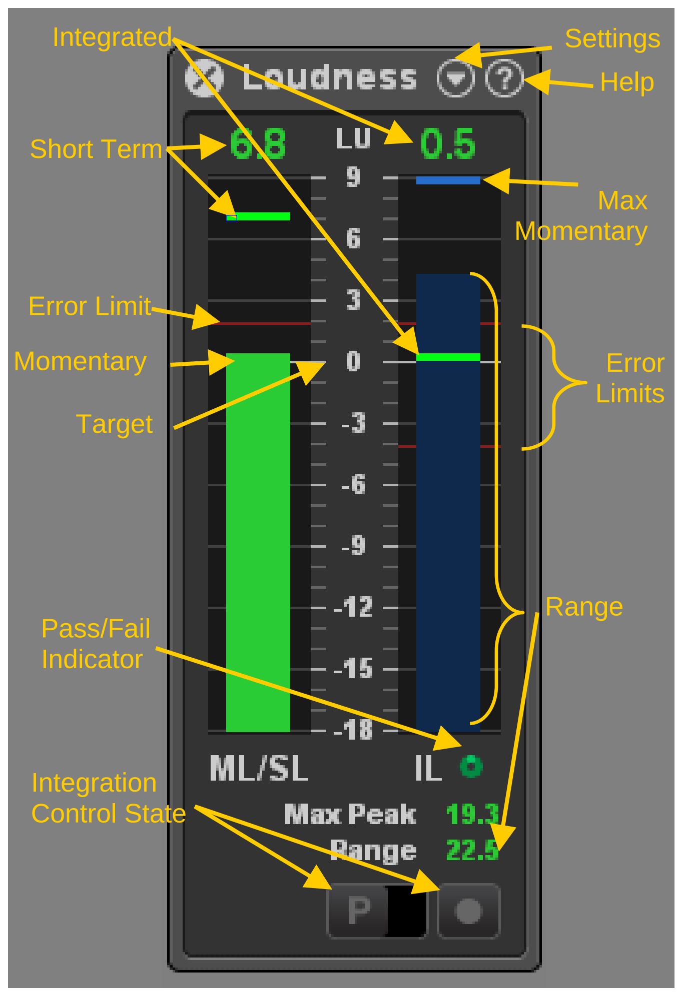

Figure 3 Loudness meter panel display

The short term loudness is shown numerically and as a horizontal line on the left bar graph while the momentary

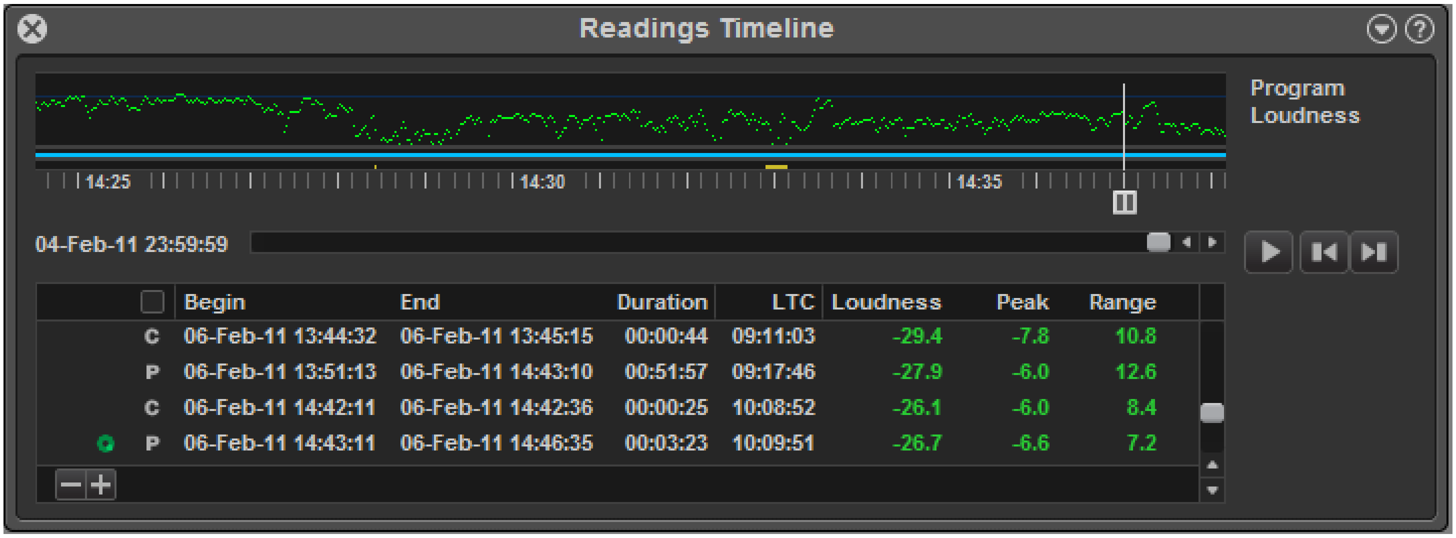

Figure 4 Readings Timeline showing the short-term loudness history and completed integrated loudness measurements of programs and commercials

loudness is shown on the bar. The momentary loudness ballistics are similar to that of a VU meter and is a very useful tool for mix engineers in monitoring program level. The short-term loudness tracks listener judgment of program loudness and is a good guide to the evolution in the integrated loudness measurement.

The second bar-graph shows the integrated loudness both in numeric and graphical form along with other parameters defined in EBU R128 including the loudness range, the maximum true peak, the maximum momentary loudness, and the maximum short term loudness. These are controlled with the same gate as the integrated loudness.

Loudness History

The short-term loudness is graphed as a function of time giving a perspective of trends, dynamics and variations. An example plot is shown in Figure 4. The cursor is used to identify any desired time in the record. All measurements corresponding to this date and time will be displayed in the UI exactly as they were when the measurements were originally made.

In addition to the loudness bargraph display, integrated loudness measurements are displayed in tabular form below the readings timeline. The left-most column shows the current measurement, identified by the swirling activity indicator. The next column identifies the measurement as a program or commercial. The next four columns identify the segment measured. The last 3 columns give the primary measurements of integrated loudness, true peak and loudness range.

Alarming

Loudness-error driven alarms can be generated by individual programs or commercials failing to meet the ATSC requirement that its loudness matches the dialnorm value. This is useful to check loudness setting operations upstream and know if the content is meeting CALM Act requirements. However, to give an “early warning” of loudness issues there is a second set of loudness alarms.

In this case, the running Short Term loudness is compared to the target loudness but using a separately specified tolerance and duration. This allows alarms to be generated when excessively loud passages occur, exceeding both an upper loudness threshold and remaining there for a specified time period. This alarm will occur even if the program or commercial segment being broadcast hasn’t completed and the resulting integrated loudness measurement made available. This gives a warning of loudness problems, enabling mitigation, before a conformance failure occurs and preventing a violation of the CALM Act.

A separate comparison is done of Short Term loudness against minimum thresholds to warn if program gets too soft. This can be applied for each type of channel (front, LFE and surrounds) independently, warning that content may be lower than listeners find desirable.

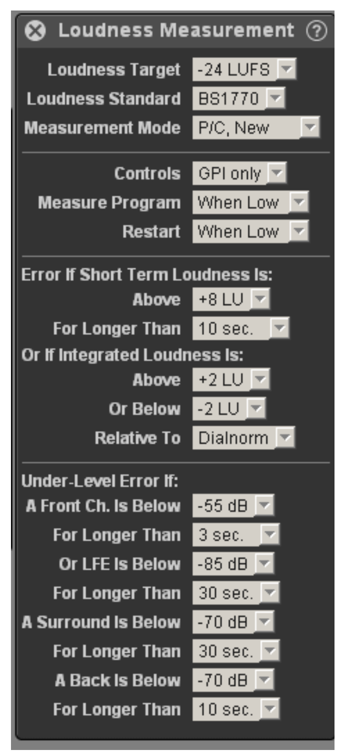

Figure 5 Loudness Settings

All loudness readings are identified in the logs as programs or commercials and stamped with date, time and time code. The readings are compared with a user specified target loudness and tolerance. A pass/fail flag is generated and if a failure occurs it can initiate user selectable alarms. Alarms always result in a visible display on the instrument’s front panel and the assertion of a hardware signal on a rear panel connector. Under user control alarms can also optionally generate audible warnings and/or email notifications. Different errors can be reported to recipients independently and multiple recipients may be specified for each error. The assertion of errors is also stored in the log files and with corresponding date, time and timecode stamp.

Forensics

FCC regulations are likely to apply only to complete segments such as an entire commercial. Strict conformance shouldn't require analysis over subsections or across multiple segments. However, the reality is that dealing with complaints will not be simple or pleasant. Even though you may escape unscathed, it is useful to understand why the viewer complained in the first place and be proactive in preventing another complaint. This requires the ability to do "forensic" analysis of the audio stream. How loud was the segment the viewer complained about? How did it compare to neighboring segments? Was it just the transition that generated the viewer unhappiness or was there something about the "dynamics" of the content? Is there a pattern to the loudness changes that trigger complaints?

These questions, and others like them, cannot be answered or solved using simplistic error reporting meters and logging schemes. The Qualis Audio Sentinel has a rich set of tools designed to make forensic investigations like these easy and effective in assuring compliance in the future. These operate through a simple cursor or tabular interface, entirely within the GUI.

In addition to saving loudness readings, the Qualis Audio logs save loudness data in sufficient detail that an accurate loudness measurement can be derived for any segment, anywhere in the audio stream. Loudness can be computed across multiple segments individually or in combination, excluding any portion for any reason. The example in Figure 6 shows the loudness, true peak and loudness range for the combination of two segments, each one minute long.

As with the real-time dedicated Loudness meter panel, the system also measures and logs the other parameters defined in EBU R128 including the loudness range, the maximum true peak, the maximum momentary loudness, and the maximum short term loudness. Each of these may be enabled for display in the GUI if desired and are automatically computed on logged data if selected. All computations are available on the stored data whether enacted at the time of measurement or during forensic examination.

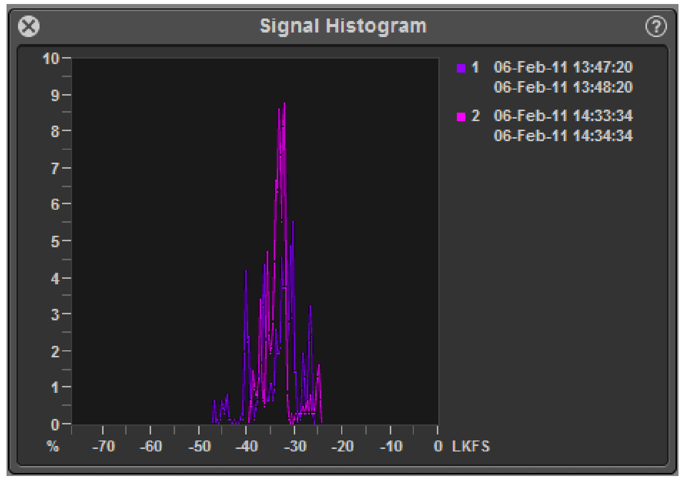

Figure 7 Loudness histogram for Figure 6

An additional tool, just one button click away, generates histograms of the segments selected on the timeline. The figure below shows the loudness distribution for the segments identified in Figure 6.

These displays show what percentage of the time the loudness was any particular value. This can give insight into the dynamic range compression used on the selected segments and help understand why a particular clip may have generated complaints. Multiple segments are plotted in different colors on the same axes, easing the comparison of one with another.

Automated Report Generation

The complete data stored in the log files allows customized report generation for quality control reporting, monitoring,

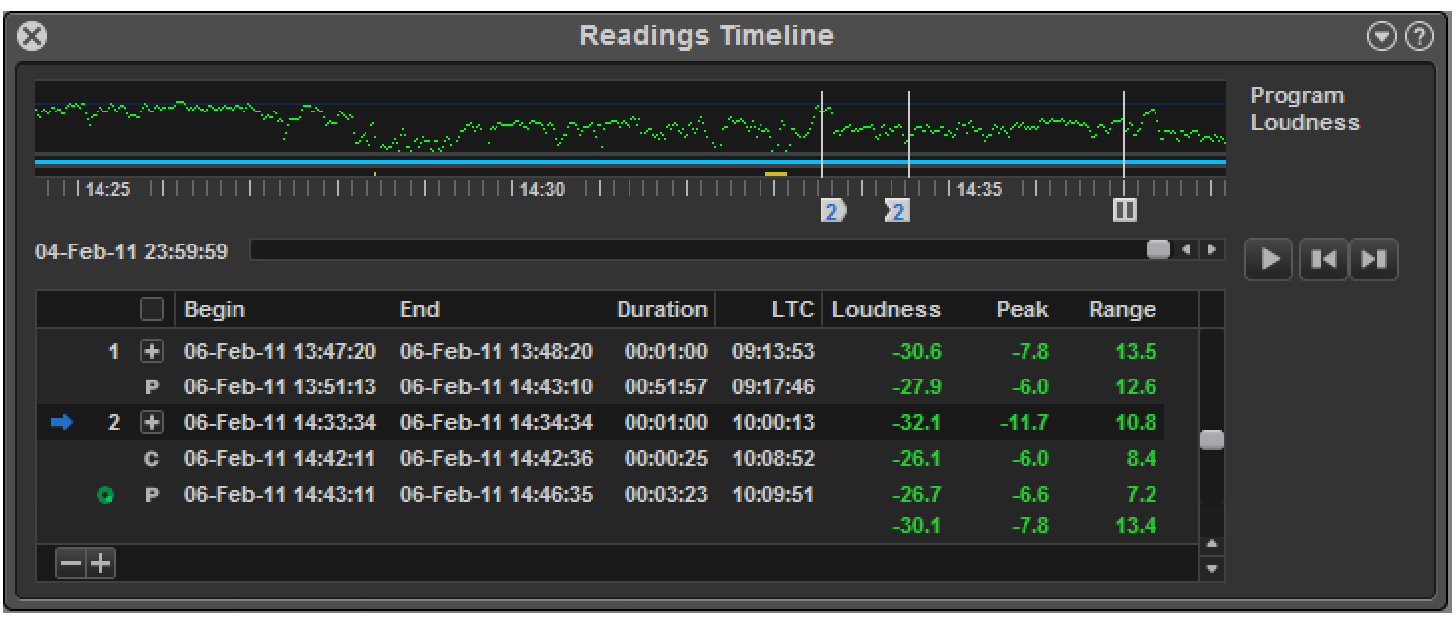

Figure 6 Readings Timeline with 2 user defined segments

query response or general management use. A time period or section of interest may be selected via the GUI cursors and saved. This data may be loaded in Excel, analyzed and displayed in whatever combination desired.

Alternately, the entire process may be run from a batch file eliminating the need for any human intervention. This approach may be used, for example, by broadcasters who want an automatically generated report each morning of the previous day's broadcasts. The script tool may load as-run logs, creating a report that lists each program and commercial along with their loudness or other desired parameters.

Improving the Loudness ROI

Implementing a loudness monitoring and logging solution provides the perfect opportunity to institute a comprehensive approach to audio quality control. In addition to covering all your loudness measurement and logging requirements, the Qualis Audio Sentinel measures virtually every parameter necessary to insure optimum audio reaches your viewers. It operates completely unattended, allowing significant reductions in labor costs compared with traditional audio QC approaches. The Sentinel solution both assures compliance with the CALM Act and provides a competitive advantage in the viewers audio experience.

Also of interest

A copy of the CALM Act may be downloaded from the Qualis Audio web site.

Tech Note #2 “Understanding and Verifying Loudness Meters” and a set of waveform files which exercise BS1770 compliant loudness meters are available on the Qualis Audio web site.

Qualis Audio, Inc.
PO Box 731
Lake Oswego, OR 97034
503.635.9376
www.qualisaudio.com
info@qualisaudio.com

We're Listening